


# Incremental Techniques for Large-Scale Dynamic Query Processing

Tutorial

**Part 1**


Iman Elghandour<sup>1</sup> Ahmet Kara<sup>2</sup> Dan Olteanu<sup>2</sup> Stijn Vansummeren<sup>1</sup>





## Dynamic query evaluation

Avoid full recomputation – compute incrementally!



## Application Scenarios:

- Real-time monitoring
- Internet of things
- Knowledge base construction
- Online machine learning

# Real-time monitoring



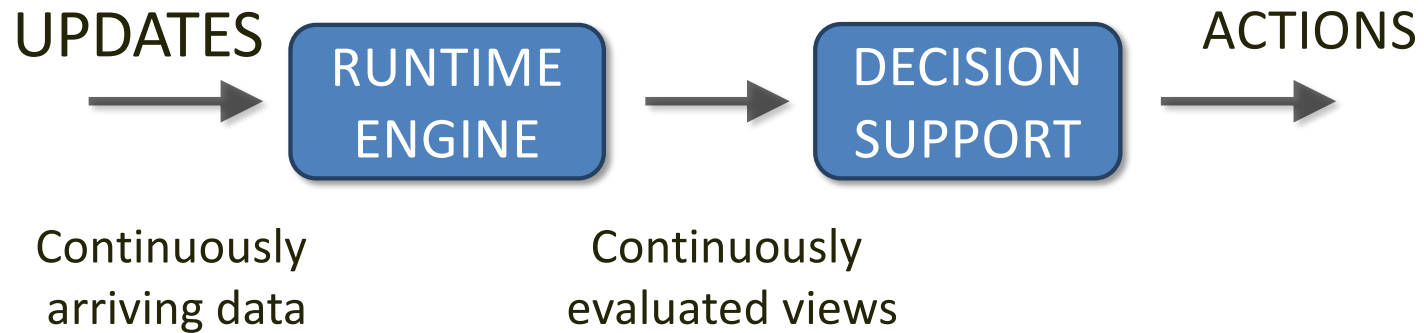
Web Analytics



Sensor Networks




Cloud Monitoring






# Complex Event Recognition



A. Artikis. Complex Event Recognition. European Conference on Artificial Intelligence (ECAI), 2016.

# Internet of things



Source: <https://blogs.cisco.com/diversity/the-internet-of-things-infographic>

# Internet of things



Source: <https://blogs.cisco.com/diversity/the-internet-of-things-infographic>

# Knowledge Base Construction

## Scalable Probabilistic Databases with Factor Graphs and MCMC

Michael Wick  
University of Massachusetts  
Computer Science  
140 Governor's Drive  
Amherst, MA  
mwick@cs.umass.edu

Andrew McCallum  
University of Massachusetts  
Computer Science  
140 Governor's Drive  
Amherst, MA  
mccallum@cs.umass.edu

Gerome Miklau  
University of Massachusetts  
Computer Science  
140 Governor's Drive  
Amherst, MA  
miklau@cs.umass.edu

Wick, McCallum, Miklau,  
PVLDB 2010.

### ABSTRACT

Incorporating probabilities into the semantics of incomplete databases has posed many challenges, forcing systems to sacrifice modeling power, scalability, or treatment of relational algebra operators. We propose an alternative approach where the underlying relational database always represents a single world, and an external factor graph encodes a distribution over possible worlds; Markov chain Monte Carlo (MCMC) inference is then used to recover this uncertainty to a desired level of fidelity. Our approach allows the efficient evaluation of arbitrary queries over probabilistic databases with arbitrary dependencies expressed by graphical models with structure that changes during inference. MCMC sampling provides efficiency by hypothesizing *modifications* to possible worlds rather than generating entire worlds from scratch. Queries are then run over the portions of the world that change, avoiding the onerous cost of running full queries over each sampled world. A significant innovation of this work is the connection between MCMC sampling and materialized view maintenance techniques: we find empirically that using view maintenance techniques is several orders of magnitude faster than naively querying each sampled world. We also demonstrate our system's ability to answer relational queries with aggregation, and demonstrate additional scalability through the use of parallelization on a real-world complex model of information extraction. This framework is sufficiently expressive to support probabilistic inference not only for answering queries, but also for inferring missing database content from raw evidence.

current PDBs do not achieve the difficult balance of expressivity and efficiency necessary to support such a range of scalable real-world structured prediction systems.

Indeed, there is an inherent tension between the expressiveness of a representation system and the efficiency of query evaluation. Many recent approaches to probabilistic databases can be characterized as residing on either pole of this continuum. For example, some systems favor efficient query evaluation by restricting modeling power with strict independence assumptions [5, 6, 1]. Other systems allow rich representations that render query evaluation intractable for a large portion of their model family [10, 24, 19, 20].

In this paper we to provide a powerful query evaluation.

Graphical models capture uncertainty and dependencies, including language processing [16], and are even more expressive when used as purpose probabilistic graphs are a part that serve as an undirected Markov random fields, and are capable of representing any exponential family probability distribution.

In our approach, we use factor graphs to represent uncertainty over our relational data, and MCMC for inference of database con-

over the portions of the world that change, avoiding the onerous cost of running full queries over each sampled world. A significant innovation of this work is the connection between MCMC sampling and materialized view maintenance techniques: we find empirically that using view maintenance techniques is several orders of magnitude faster than naively querying each sampled world. We

# Knowledge Base Construction

## Incremental Knowledge Base Construction Using DeepDive

Shin et al,  
PVLDB 2015.

Jaeho Shin<sup>†</sup> Sen Wu<sup>†</sup> Feiran Wang<sup>†</sup> Christopher De Sa<sup>†</sup> Ce Zhang<sup>†‡</sup> Christopher Ré<sup>†</sup>

<sup>†</sup>Stanford University

<sup>‡</sup>University of Wisconsin-Madison

{jaeho, senwu, feiran, cdesa, czhang, chrismre}@cs.stanford.edu

### ABSTRACT

Populating a database with unstructured information is a long-standing problem in industry and research that encompasses problems of extraction, cleaning, and integration. Recent names used for this problem include dealing with dark data and knowledge base construction (KBC). In this work, we describe DeepDive, a system that combines database and machine learning ideas to help develop KBC systems, and we present techniques to make the KBC process more efficient.

We observe that the KBC process is iterative, and we develop techniques to incrementally produce inference results for KBC systems. We propose two methods for incremental inference, based respectively on sampling and variational techniques. We also study the tradeoff space of these methods and develop a simple rule-based optimizer. DeepDive includes all of these contributions, and we evaluate DeepDive on five KBC systems, showing that it can speed up KBC inference tasks by up to two orders of magnitude with negligible impact on quality.

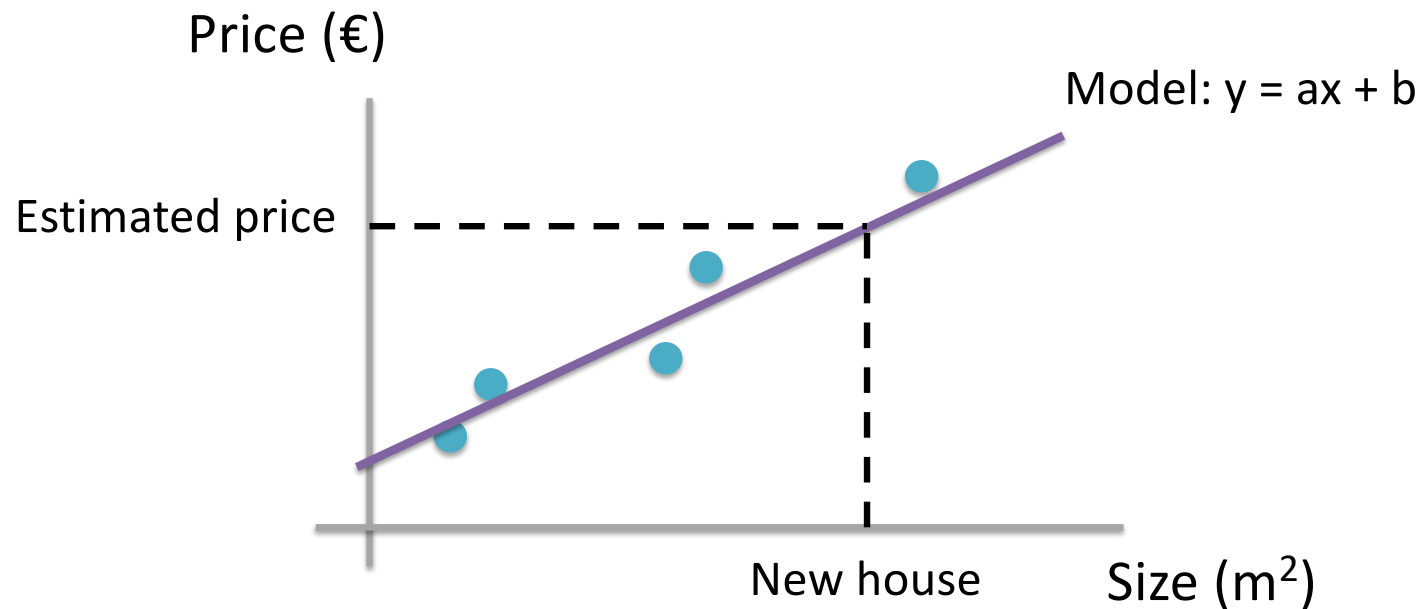
complex relationships. Typically, quality is assessed using two complementary measures: precision (how often a claimed tuple is correct) and recall (of the possible tuples to extract, how many are actually extracted). These systems can ingest massive numbers of documents—far outstripping the document counts of even well-funded human curation efforts. Industrially, KBC systems are constructed by skilled engineers in a months-long process. In such systems, this question spans including program focus on a narrower rapid the program loop, the more quick.

This paper presents knowledge base construction.<sup>1</sup> DeepDive’s language and execution model are similar to other KBC systems: DeepDive uses a high-level declarative language [11, 28, 30]. From a

We observe that the KBC process is iterative, and we develop techniques to incrementally produce inference results for KBC systems. We propose two methods for incremental inference, based respectively on sampling and variational techniques. We also study the tradeoff space of these meth-

# Online Machine Learning

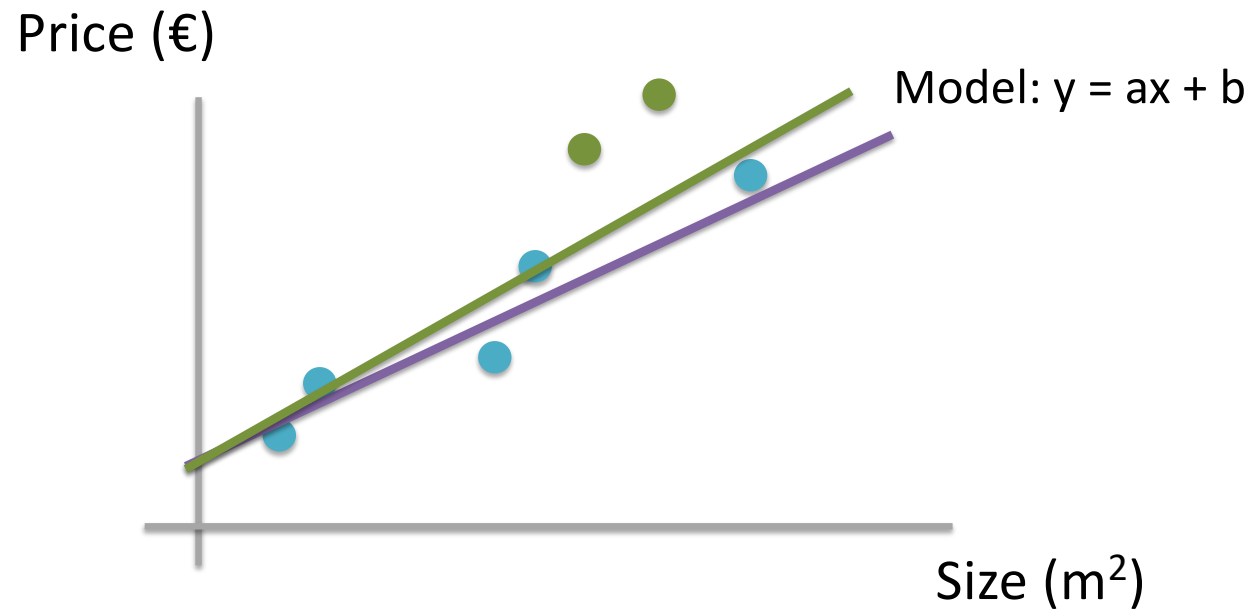
## Estimating House Prices



Linear regression with parameters  $(a,b)$

# Online Machine Learning

## Estimating House Prices



Linear regression with parameters  $(a,b)$



## Conclusion:

Dynamic Query Processing is pervasive in a wide range of application areas.




# Outline

- Part I: Introduction
- Part II: Main Algorithmic Ideas in Dynamic Query Processing: Traditional IVM and Recent Advances
- Part III: Generalizations to Arbitrary Ring Structures
- Part IV: Dynamic Query Processing in Big Data Frameworks
- Part V: Outlook


# Outline

- Part I: Introduction
- ➔ • **Part II: Main Algorithmic Ideas in Dynamic Query Processing: Traditional IVM and Recent Advances**
- Part III: Generalizations to Arbitrary Ring Structures
- Part IV: Dynamic Query Processing in Big Data Frameworks
- Part V: Outlook



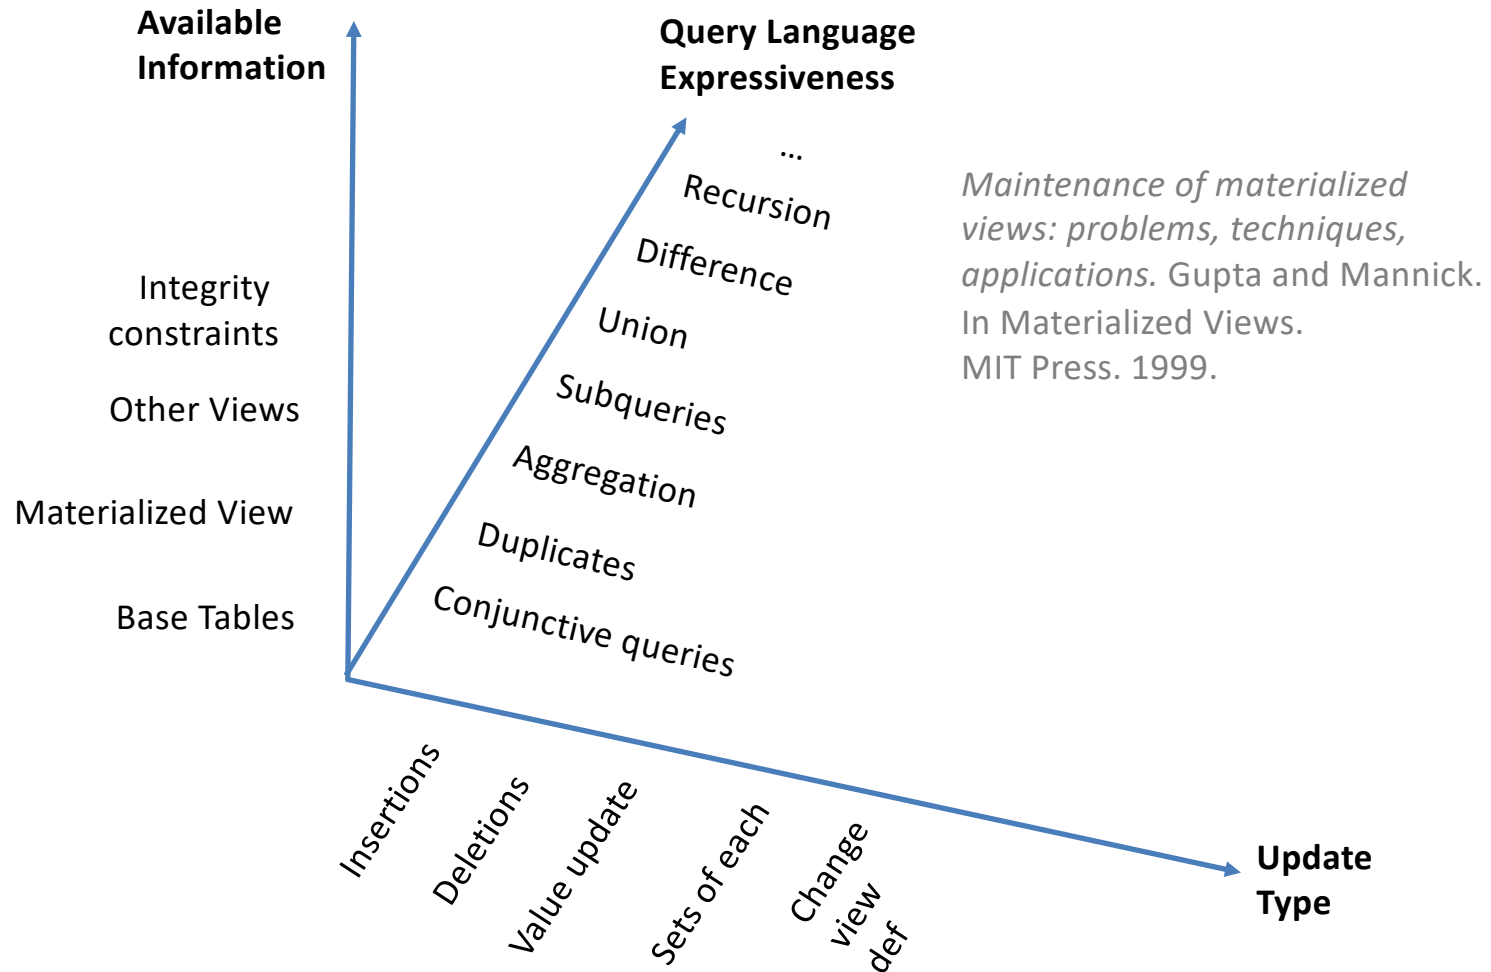
## Dynamic query evaluation

Avoid full recomputation – compute incrementally

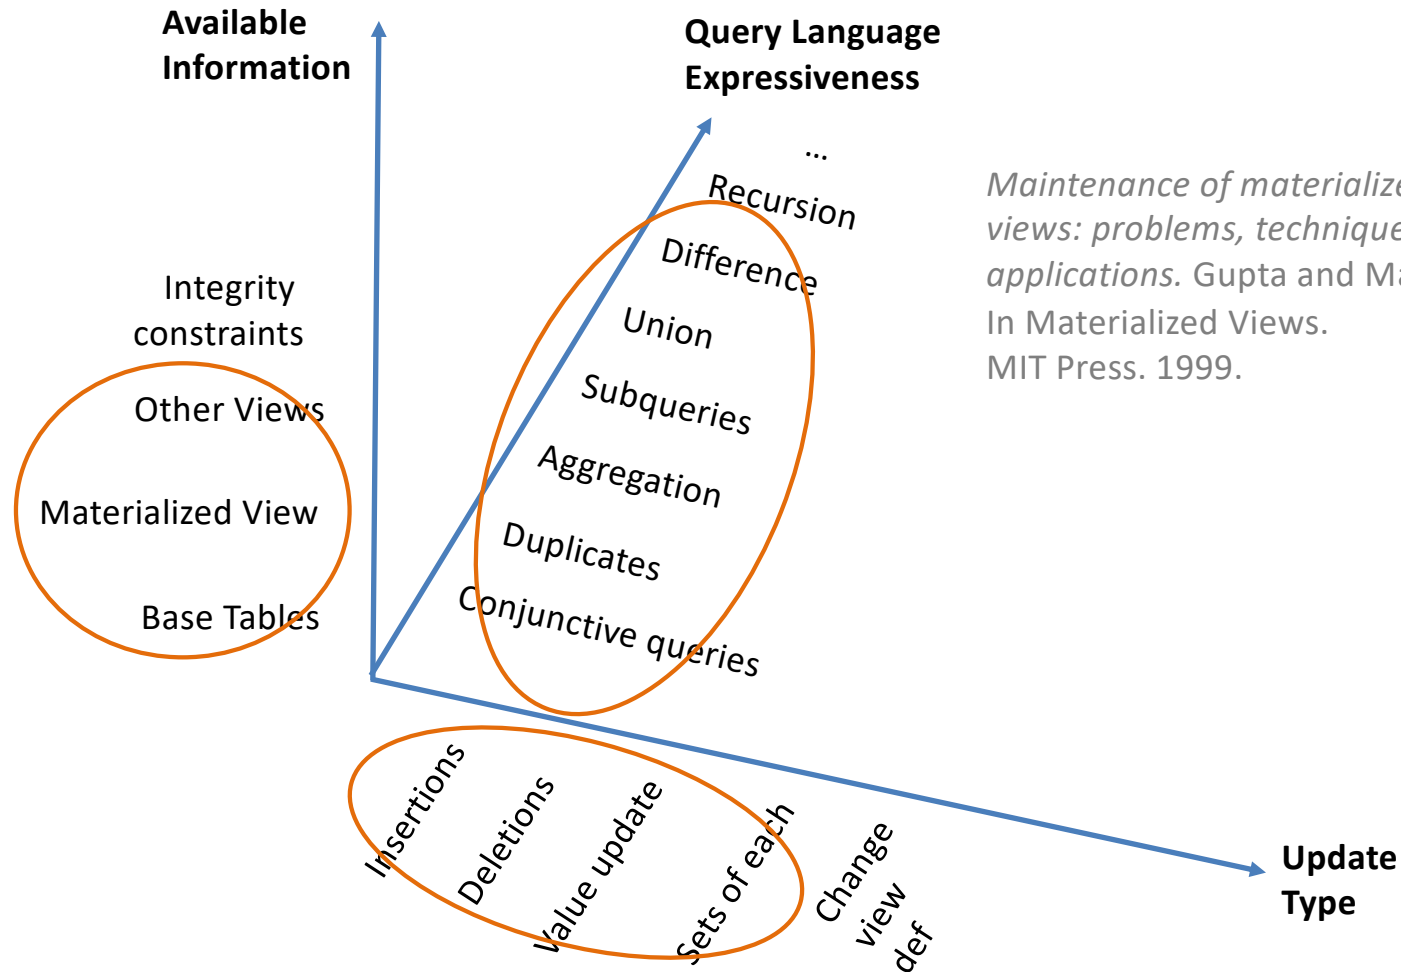


## Incremental View Maintenance (IVM)

# Dimensions of IVM




# Dimensions of IVM: this tutorial



*Maintenance of materialized views: problems, techniques, applications.* Gupta and Mannick. In *Materialized Views*. MIT Press. 1999.

# Main Algorithmic Ideas

1. IVM  $\equiv$  processing of delta queries
2. Materialize results of subqueries in addition to the actual query result
3. Exploit data skew



# Main Algorithmic Ideas

➔ 1. IVM  $\equiv$  processing of delta queries

2. Materialize results of subqueries in addition to the actual query result

3. Exploit data skew

Time

1993

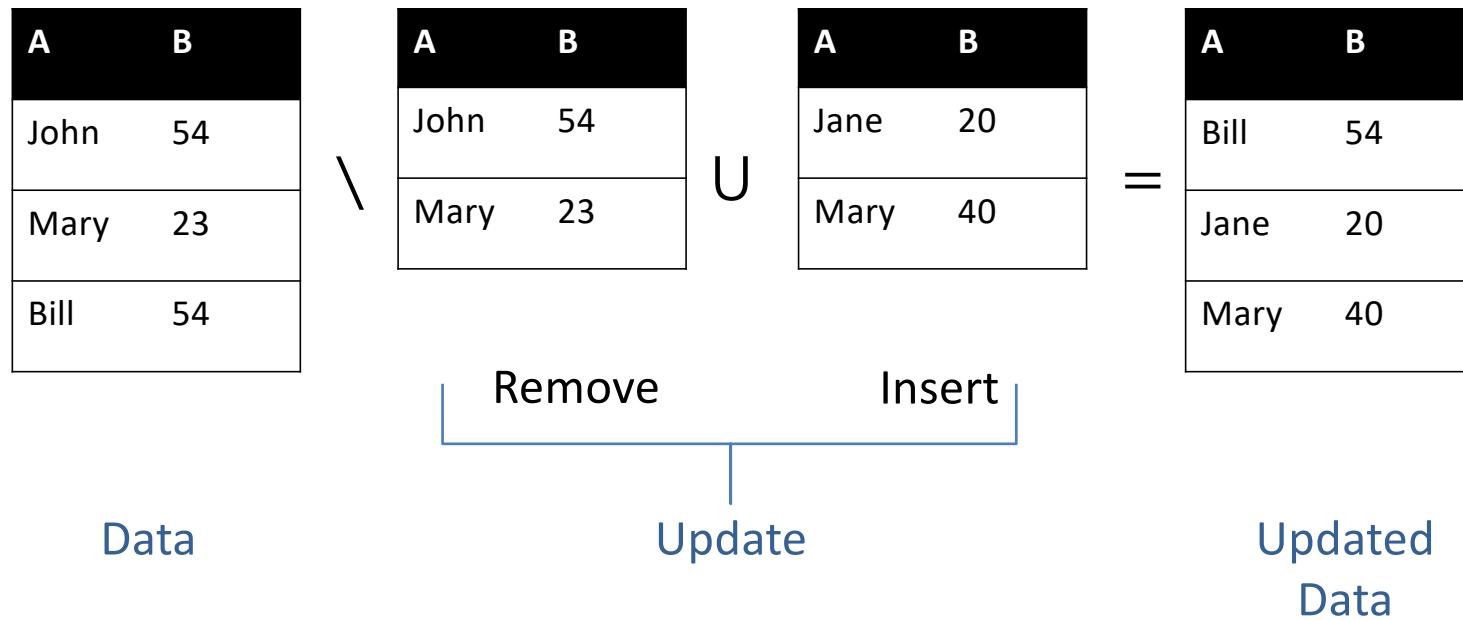
1997

2009

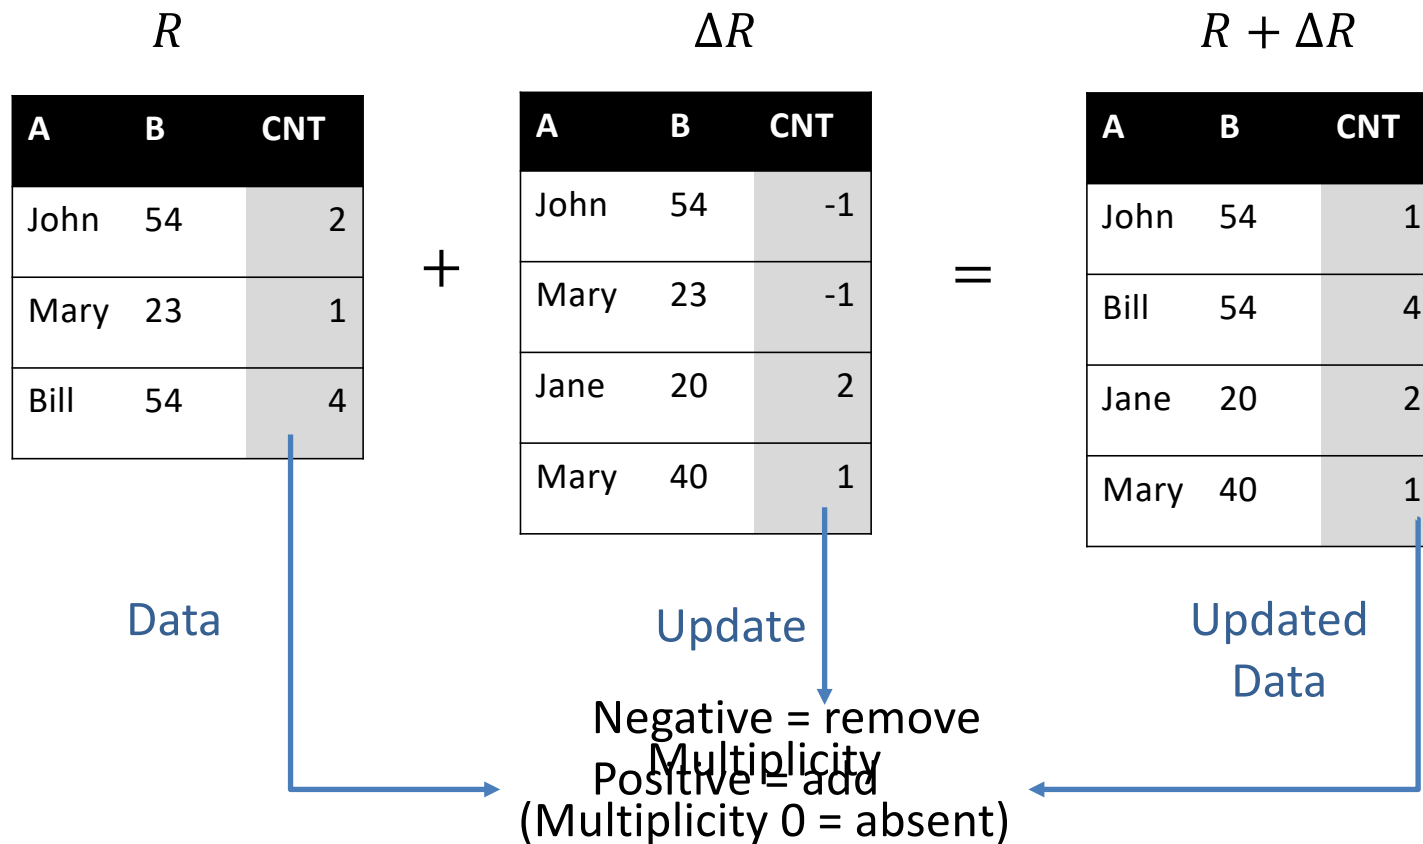
2018



# Traditional update representation



# Uniform update representation



# Query semantics (1/5)

## Selection

$\sigma_{B>30}$

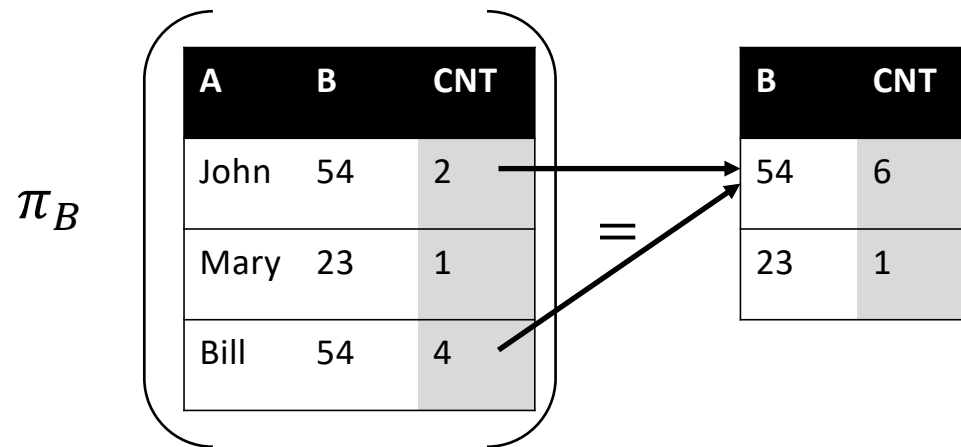
A	B	CNT
John	54	2
Mary	23	1
Bill	54	4

=

A	B	CNT
John	54	2
Bill	54	4

## Query semantics (2/5)


### Projection



Duplicate-preserving bag-based projection

## Query semantics (2/5)

### Projection



Duplicate-preserving bag-based projection

## Query semantics (3/5)

### Union

A	B	CNT
John	54	2
Mary	23	1
Bill	54	4

 $+$ 

A	B	CNT
John	54	-1
Mary	23	-1
Jane	20	2
Mary	40	1

 $=$ 

A	B	CNT
John	54	1
Bill	54	4
Jane	20	2
Mary	40	1

Duplicate-preserving bag-based union

## Query semantics (4/5)

### Difference

A	B	CNT
John	54	2
Mary	23	1
Bill	54	4

—

A	B	CNT
John	54	-1
Mary	23	-1
Jane	20	2
Mary	40	1

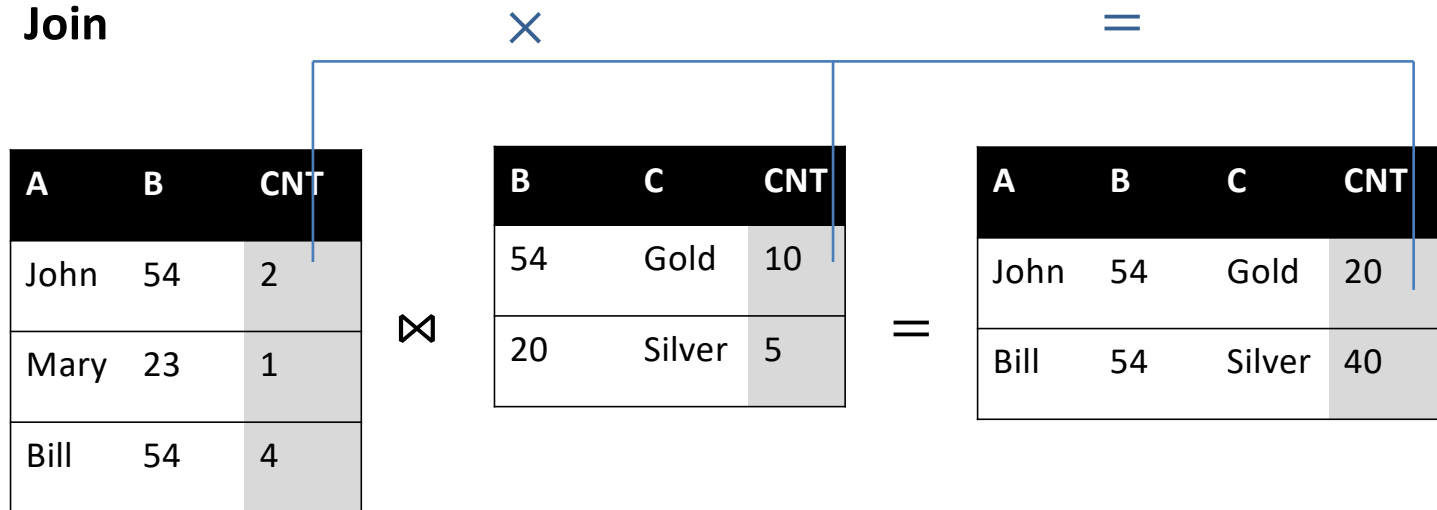
=

A	B	CNT
John	54	3
Mary	23	2
Bill	54	4
Jane	20	-2
Mary	40	-1

This is *\*not\** bag difference!

## Query semantics (5/5)

Join



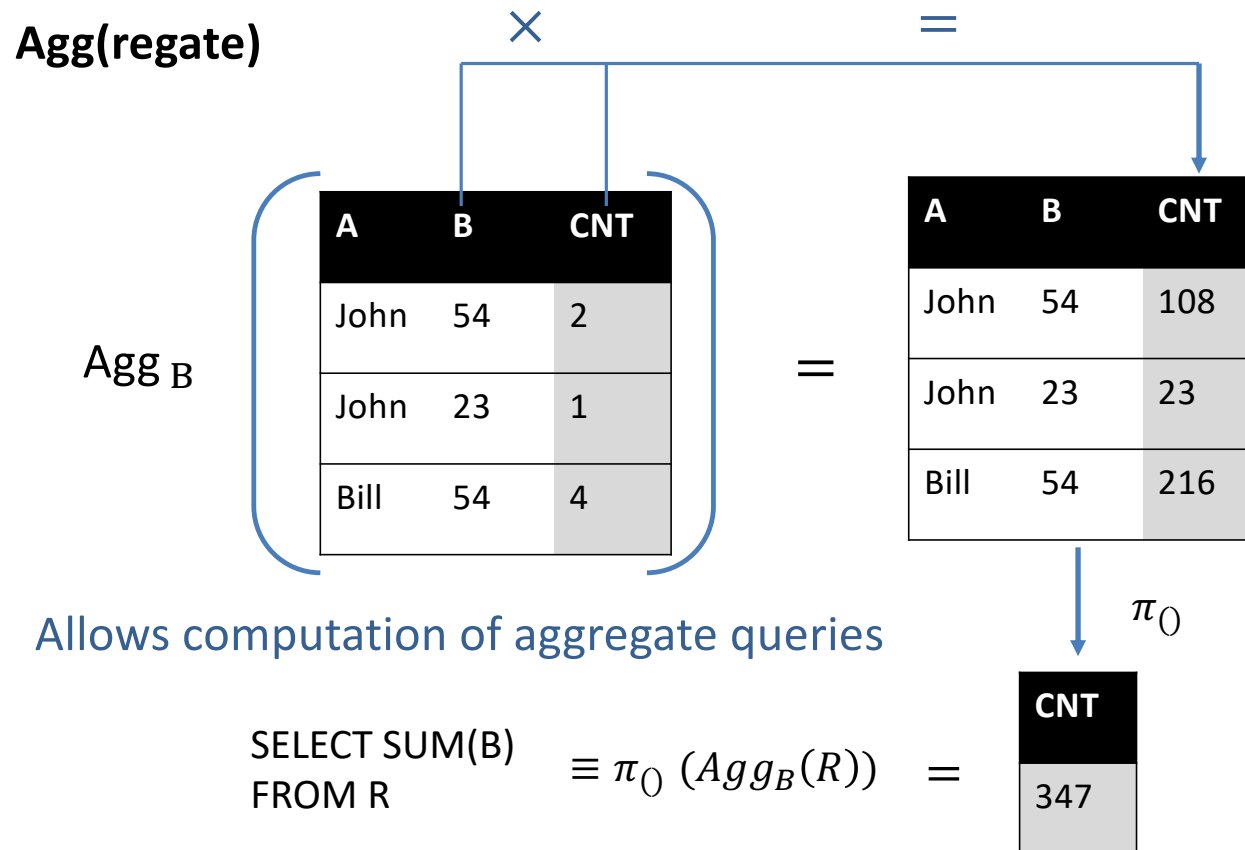
Multiply multiplicities of joining tuples



## Observations

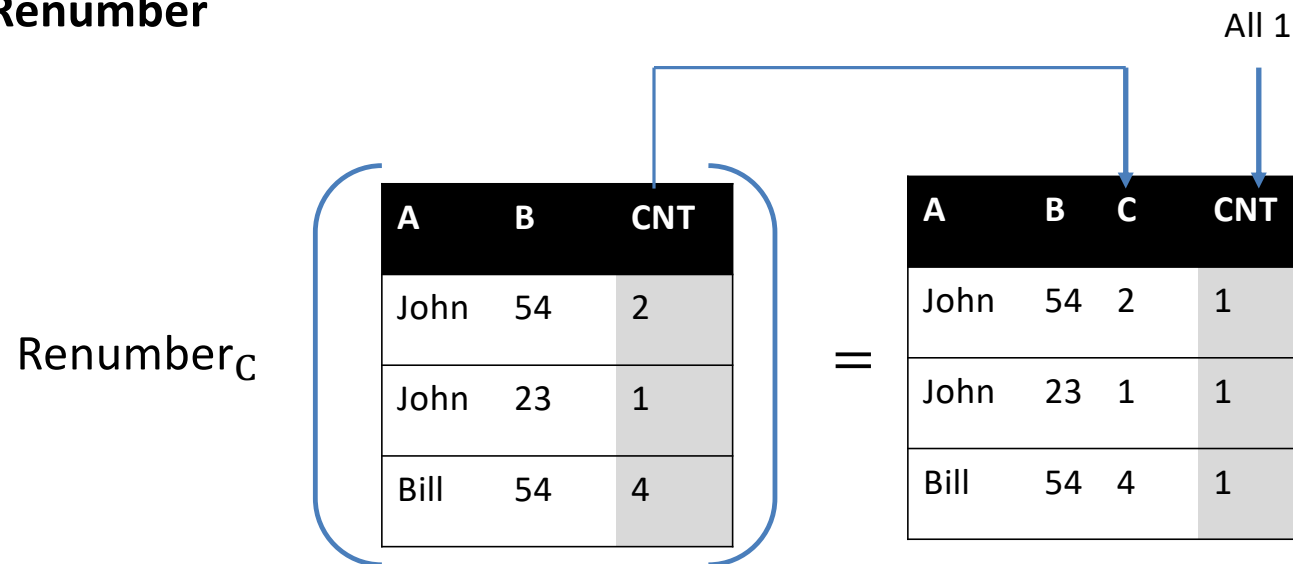
- Query evaluation algorithms are trivially modified to compute the CNT values.
- Under this modified semantics, each tuple in the query result specifies the number of derivations for that tuple.

# Query semantics: aggregation



# Query semantics: aggregation

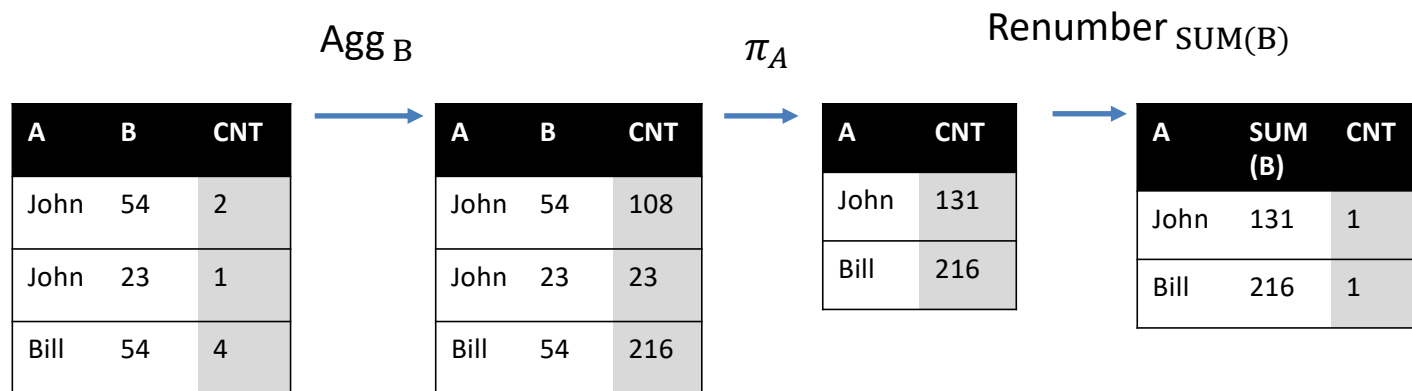
## Renumber



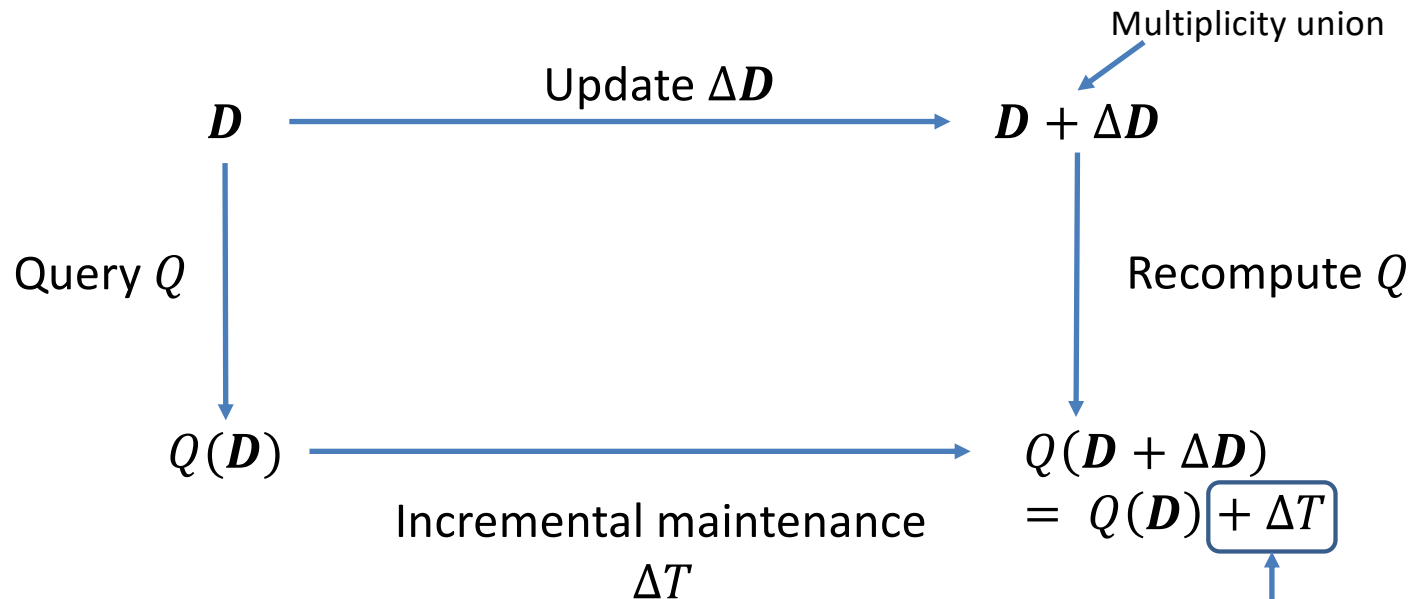
Allows computation of groupby + aggregate queries

# Groupby: Example

SELECT SUM(B)  
FROM R  
GROUP BY A



# Delta queries



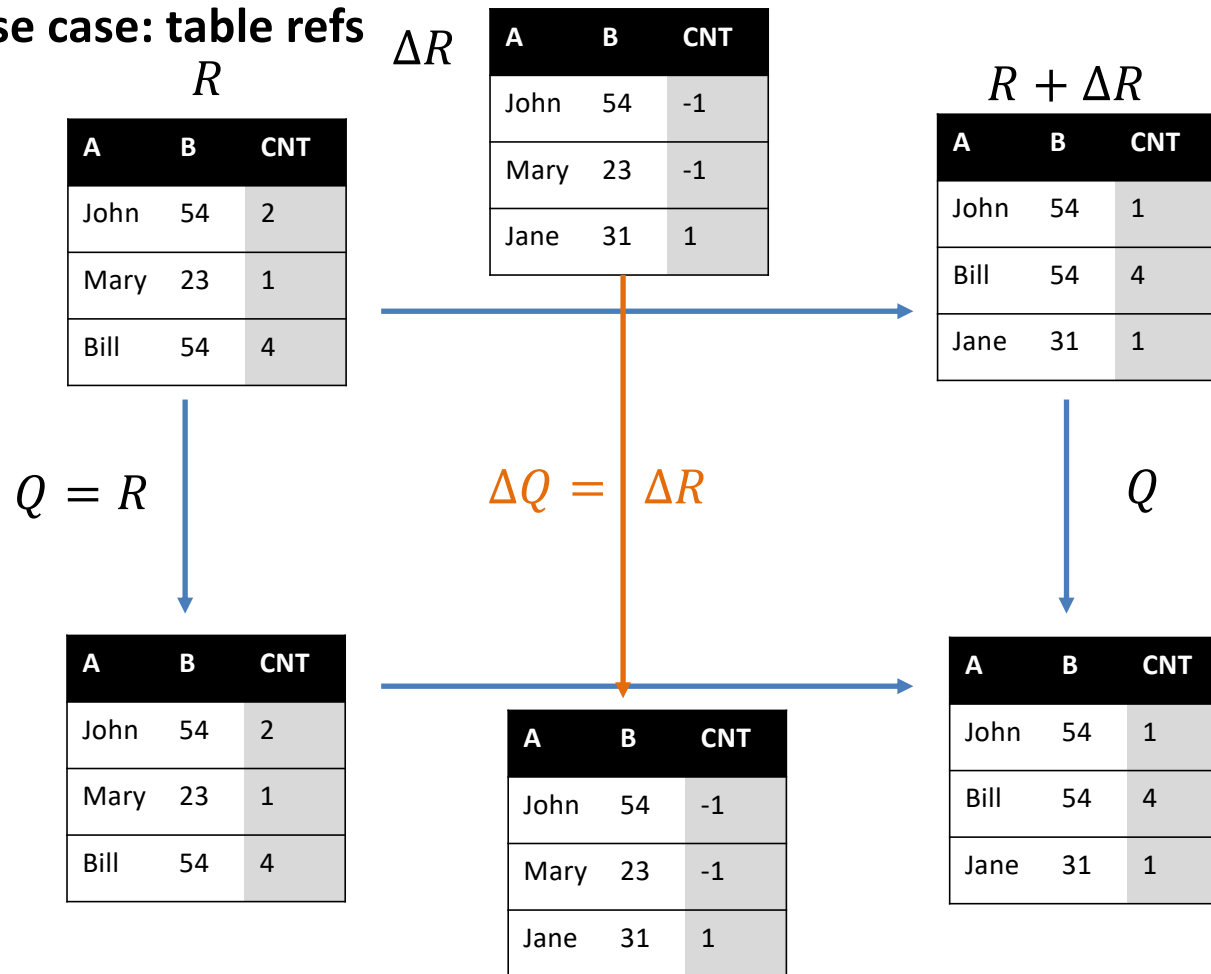
## Key insight:

For every query  $Q$  in relational algebra (with multiplicities), it is possible to write a query  $\Delta Q$  that operates on the old database  $D$  and the update  $\Delta D$  s.t.

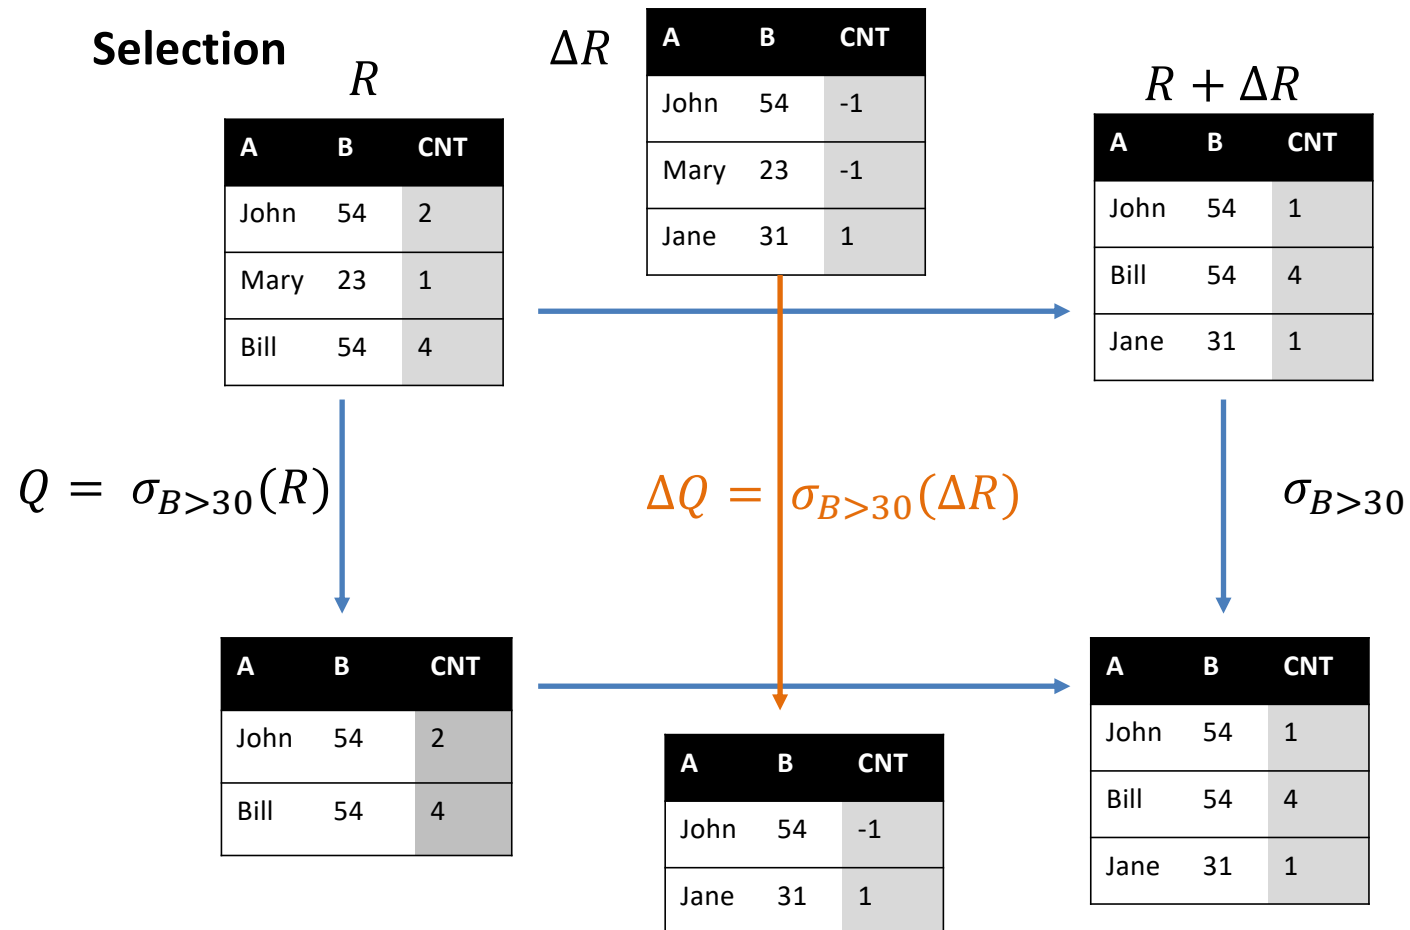
$$Q(D + \Delta D) = Q(D) + \Delta Q(D, \Delta D)$$

# Delta queries

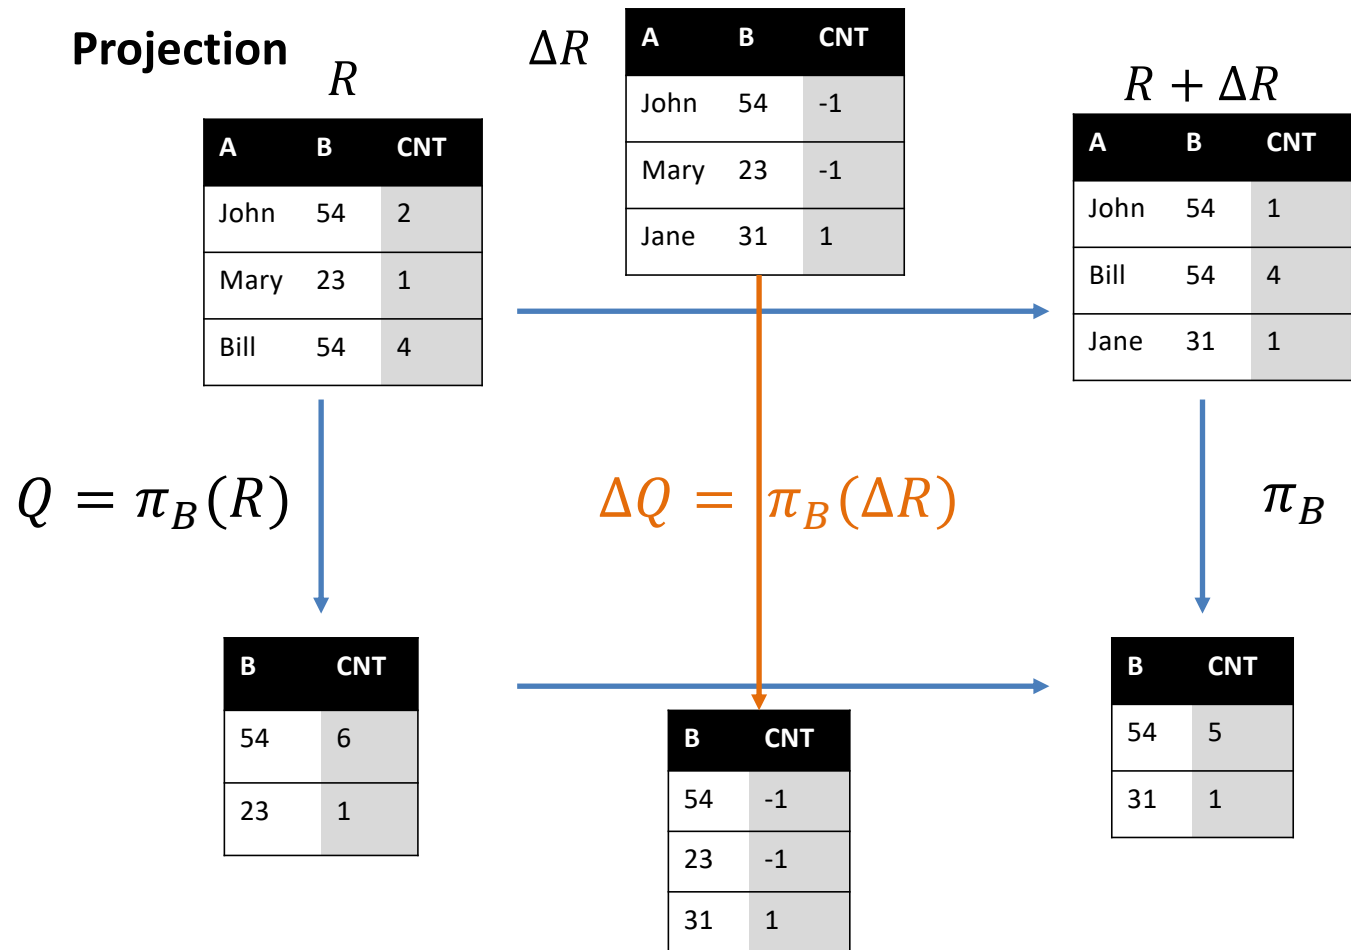
Base case: table refs



# Delta queries



# Delta queries

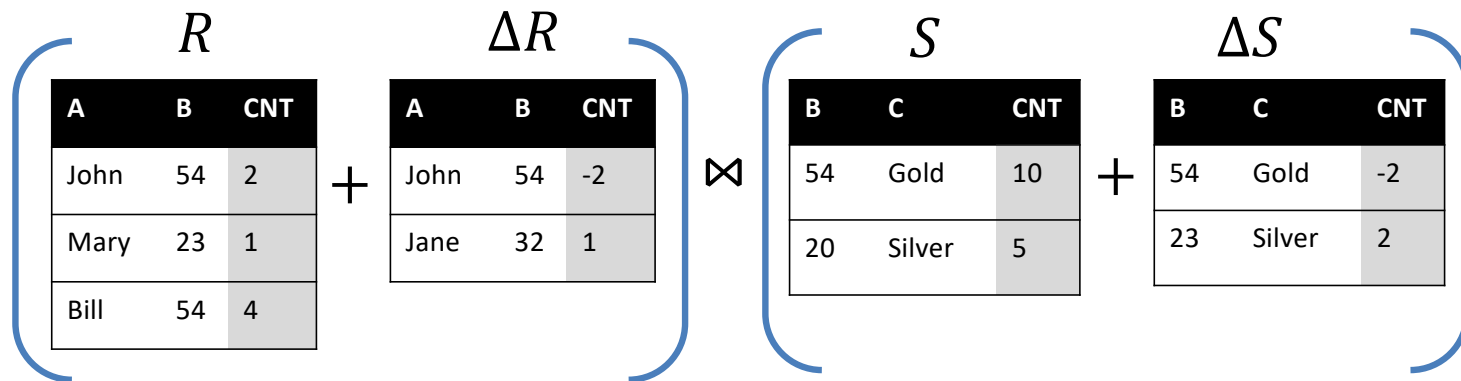




# Delta queries

Join

$$Q = R \bowtie S$$



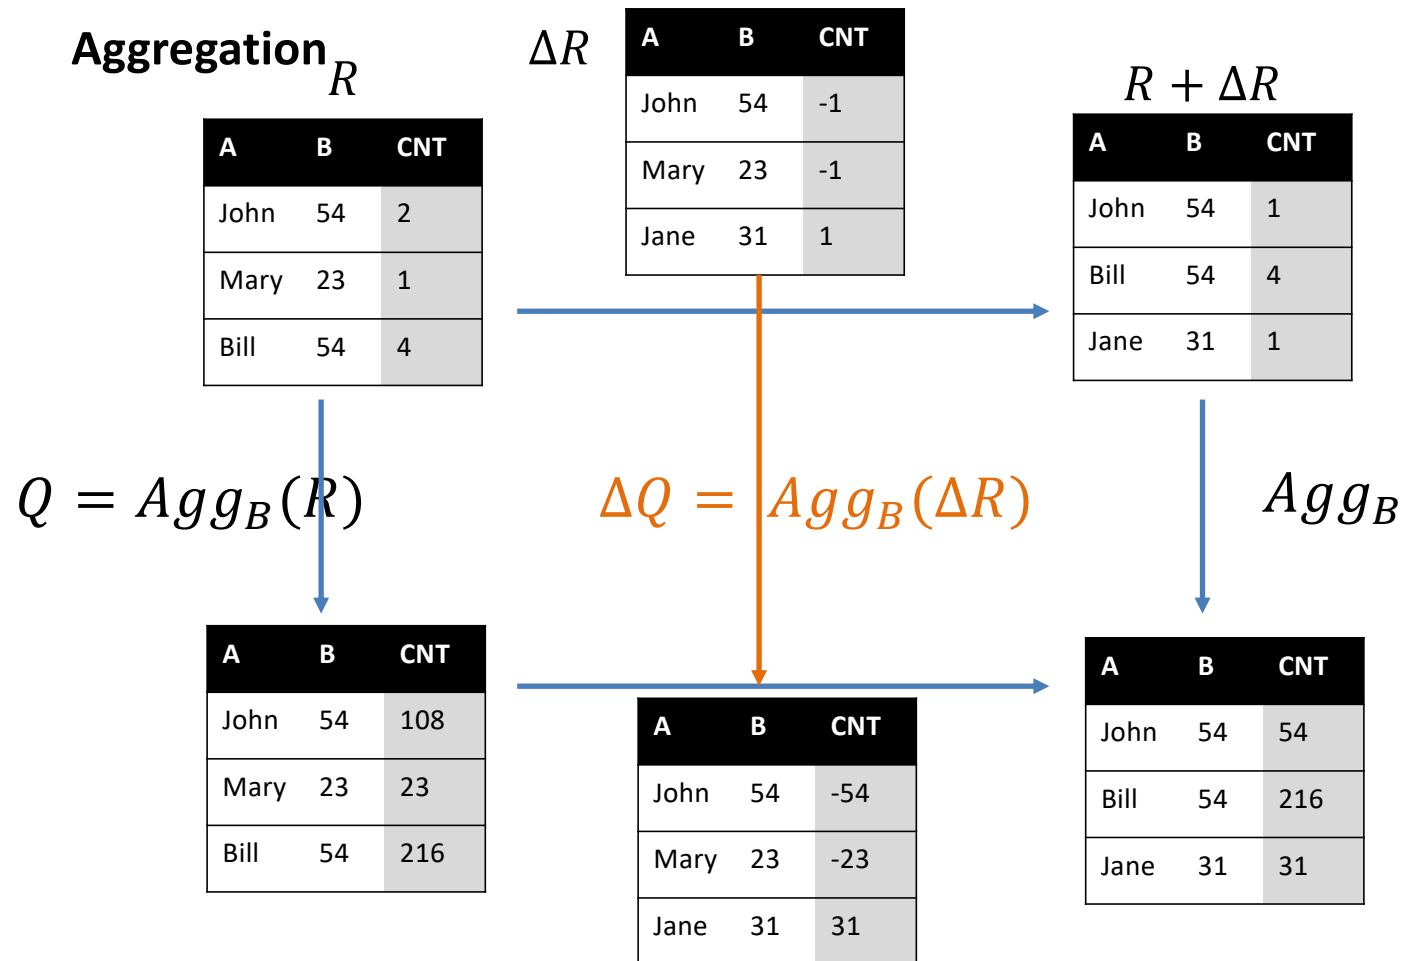
$$Q(D + \Delta D) = (R + \Delta R) \bowtie (S + \Delta S) \quad \text{Distributivity of } + \text{ over } \bowtie$$

$$= (R \bowtie S) + (\Delta R \bowtie S) + (R \bowtie \Delta S) + (\Delta R \bowtie \Delta S)$$

$$Q(D)$$

$$\Delta Q(D, \Delta D)$$

# Delta queries



# Delta queries

**Renumber**       $Q = \text{Renumber}_C(R)$

$Q$   $\left( \begin{array}{c} R \\ \Delta R \end{array} \right) + =$

A	B	CNT
John	54	2
John	23	1
Bill	54	4

A	B	CNT
John	54	-2
Jane	32	1

A	B	C	CNT
John	23	1	1
Bill	54	4	1
Jane	32	1	1

$\Delta \text{Renumber}_C(R, \Delta R)$

$= \text{Renumber}_C(R + \Delta R) - \text{Renumber}_C(R)$

**Recomputation necessary!**  
**Specialized algorithm possible**

## Delta queries: summary


Query $Q(D)$	Delta Query $\Delta Q(D, \Delta D)$
Table $R$	Table $\Delta R$
$\sigma_{\theta}(Q')$	$\sigma_{\theta}(\Delta Q')$
$\pi_{\vec{A}}(Q')$	$\pi_{\vec{A}}(\Delta Q')$
$Q_1 + Q_2$	$\Delta Q_1 + \Delta Q_2$
$Q_1 \bowtie Q_2$	$\Delta Q_1 \bowtie Q_2 + Q_1 \bowtie \Delta Q_2 + \Delta Q_1 \bowtie \Delta Q_2$
$\text{Agg}_A(Q')$	$\text{Agg}_A(\Delta Q')$
$\text{Renumber}_A(Q')$	$\text{Renumber}_A(Q' + \Delta Q') - \text{Renumber}_A(Q')$

# The Counting Algorithm

*Maintaining Views Incrementally.*

Gupta, Mumick, Subrahmaniam. SIGMOD 1993.

- Store all relations in database  $D$
- Store (materialize)  $Q(D)$  in view  $V$
- Upon update  $\Delta D$ :
  - Use  $\Delta Q$  to compute  $\Delta Q(D, \Delta D)$
  - Add  $\Delta Q(D, \Delta D)$  to  $V$



Use your favorite Query Evaluation algorithm to evaluate this.  $\Delta D$  is expected to be small!

# The Counting Algorithm: an example

$$Q = \pi_{AD}(R \bowtie S \bowtie T)$$

$R$			$S$			$T$			$Q(D)$		
A	B	CNT	B	C	CNT	C	D	CNT	A	D	CNT
John	54	2	54	Gold	2	Gold	100	2	John	100	8
Mary	23	1	20	Silver	1	Gold	80	1	John	80	4
Bill	54	4				Silver	50	1	Bill	100	16
						Bronze	20	4	Bill	80	8

$\Delta R$		
A	B	CNT
John	54	-2
Mary	23	2

$$\Delta Q = \pi_{AD}(\Delta R \bowtie S \bowtie T$$

$$+ (R + \Delta R) \bowtie \Delta S \bowtie T$$

$$+ (R + \Delta R) \bowtie (S + \Delta S) \bowtie \Delta T)$$
  

$\Delta Q(D, \Delta D)$		
A	D	CNT
John	100	-8
John	80	-4

Empty!

# The Counting Algorithm: an example

$$Q = \pi_{AD}(R \bowtie S \bowtie T)$$

*R*

A	B	CNT
John	54	2
Mary	23	1
Bill	54	4

*S*

B	C	CNT
54	Gold	2
20	Silver	1

*T*

C	D	CNT
Gold	100	2
Gold	80	1
Silver	50	1
Bronze	20	4

*Q(D)*

A	D	CNT
John	100	8
John	80	4
Bill	100	16
Bill	80	8

$\Delta R$

A	B	CNT
John	54	-2
Mary	23	2

$\Delta Q(D, \Delta D)$

A	D	CNT
John	100	-8
John	80	-4

$$\Delta_R Q = \pi_{AD}(\Delta R \bowtie S \bowtie T)$$

# Main Algorithmic Ideas

1. IVM  $\equiv$  processing of delta queries
- ➔ 2. Materialize results of subqueries in addition to the actual query result
3. Exploit data skew






# Delta evaluation through recomputation?

$$Q = \pi_{AD}(R \bowtie S \bowtie T)$$

Maintain under updates to R



$$\Delta_R Q = \pi_{AD}(\Delta R \bowtie S \bowtie T)$$



- The join  $S \bowtie T$  is recomputed for every update  $\Delta R$
- Update latency may be high

# Key Insight

$$Q = \pi_{AD}(R \bowtie S \bowtie T)$$

$$\equiv \pi_{AD}(R \bowtie \pi_{BD}(S \bowtie T))$$

$R$   $S$

Maintain under updates to R

A	B	CNT
John	54	2
Mary	23	1
Bill	54	4

B	C	CNT
54	Gold	2
20	Silver	1

C	D	CNT
Gold	100	2
Gold	80	1
Silver	50	1
Bronze	20	4

$Q(D)$

A	D	CNT
John	100	8
John	80	4
Bill	100	16
Bill	80	8

$\Delta R$

A	B	CNT
John	54	-2
Mary	23	2

$\pi_{BD}(S \bowtie T)$

B	D	CNT
54	100	4
54	80	2
20	Silver	1

Maintain under  
updates to  
S and T

$\Delta Q(D, \Delta D)$

A	D	CNT
John	100	-8
John	80	-4

$$\Delta_R Q = \pi_{AD}(\Delta R \bowtie S \bowtie T)$$

$$\equiv \pi_{AD}(\Delta R \bowtie \pi_{BD}(S \bowtie T))$$

$Q_{ST}$

Optimize  
Materialize as auxiliary view  
to avoid recomputation

# Key Insight

$$Q_{ST} = \pi_{BD}(S \bowtie T)$$

How to maintain  $Q_{ST}$

$R$			$S$			$T$			$\pi_{BD}(S \bowtie T)$		
A	B	CNT	B	C	CNT	C	D	CNT	B	D	CNT
John	54	2	54	Gold	2	Gold	100	2	54	100	4
Mary	23	1	20	Silver	1	Gold	80	1	54	80	2
Bill	54	4				Silver	50	1	20	Silver	1
						Bronze	20	4			

$$\Delta_S Q_{ST} = \pi_{BD}(\Delta S \bowtie T)$$

$$\Delta_T Q_{ST} = \pi_{BD}(S \bowtie \Delta T)$$

No auxiliary view necessary (base table)

Trivial to maintain under updates

(When really a subquery: continue reasoning)

## What are we doing here ?

$$Q = \pi_{AD}(R \bowtie S \bowtie T)$$

Maintain under updates to R

$$\Delta_R Q = \pi_{AD}(\Delta R \bowtie \pi_{BD}(S \bowtie T))$$

Materialize as aux. view

$$Q_{ST} = \pi_{BD}(S \bowtie T)$$

Maintain under updates to S, T

$$\Delta_S Q_{ST} = \pi_{BD}(\Delta S \bowtie T)$$

Materialize T

$$\Delta_T Q_{ST} = \pi_{BD}(S \bowtie \Delta T)$$

Materialize S

First-order Delta Query

## What are we doing here ?

$$Q = \pi_{AD}(R \bowtie S \bowtie T)$$

Maintain under updates to R

$$\Delta_R Q = \pi_{AD}(\Delta R \bowtie \pi_{BD}(S \bowtie T))$$

Materialize as aux. view

$$Q_{ST} = \pi_{BD}(S \bowtie T)$$

Maintain under updates to S, T

$$\Delta_S Q_{ST} = \pi_{BD}(\Delta S \bowtie T)$$

Materialize T

$$\Delta_T Q_{ST} = \pi_{BD}(S \bowtie \Delta T)$$

Materialize S

Delta of (subquery of)  
a Delta

≡

**Higher-Order Delta**

# Continuing our reasoning

$$Q = \pi_{AD}(R \bowtie S \bowtie T)$$

$$\equiv \pi_{AD}(\pi_{AC}(R \bowtie S) \bowtie T)$$

$R$                        $S$

A	B	CNT
John	54	2
Mary	23	1
Bill	54	4

B	C	CNT
54	Gold	2
20	Silver	1

C	D	CNT
Gold	100	2
Gold	80	1
Silver	50	1
Bronze	20	4

$Q(D)$

A	D	CNT
John	100	8
John	80	4
Bill	100	16
Bill	80	8

$\pi_{AC}(R \bowtie S)$

A	C	CNT
John	Gold	4
Bill	Gold	8

Maintain under  
updates to  
R and S

$\Delta T$

C	D	CNT
Gold	100	-1
Bronze	20	2

$\Delta_T Q(D, \Delta D)$

A	D	CNT
John	100	-4
John	100	-8

$$\Delta_T Q = \pi_{AD}(R \bowtie S \bowtie \Delta T)$$

$$\equiv \pi_{AD}(\pi_{AC}(R \bowtie S) \bowtie \Delta T)$$

Materialize as auxiliary view  
to avoid recomputation

## Continuing our reasoning

$$Q = \pi_{AD}(R \bowtie S \bowtie T)$$

$$\equiv \pi_{AD}(R \bowtie T \bowtie S)$$

$R$   $S$

A	B	CNT
John	54	2
Mary	23	1
Bill	54	4

B	C	CNT
54	Gold	2
20	Silver	1

$\Delta S$

B	C	CNT
54	Gold	3
23	Silver	1

Maintain under updates to S

$T$

C	D	CNT
Gold	100	2
Gold	80	1
Silver	50	1
Bronze	20	4

$Q(D)$

A	D	CNT
John	100	8
John	80	4
Bill	100	16
Bill	80	8

$\Delta Q(D, \Delta D)$

A	D	CNT
Mary	50	1

$$\Delta_T Q = \pi_{AD}(R \bowtie \Delta S \bowtie T)$$

$$\equiv \pi_{AD}(\boxed{R \bowtie T} \bowtie \Delta S)$$

Could materialize this, but it is a Cartesian product; doesn't perform better than re-evaluation

# Key Insight: conclusion

$$Q = \pi_{AD}(R \bowtie S \bowtie T)$$

Maintain under all updates

Maintain base tables + query result ...

$R$			$S$			$T$			$Q(D)$		
A	B	CNT	B	C	CNT	C	D	CNT	A	D	CNT
John	54	2	54	Gold	2	Gold	100	2	John	100	8
Mary	23	1	20	Silver	1	Gold	80	1	John	80	4
Bill	54	4				Silver	50	1	Bill	100	16
						Bronze	20	4	Bill	80	8

... as well as query subresults as auxiliary views

$\pi_{AC}(R \bowtie S)$			$\pi_{BD}(S \bowtie T)$		
A	C	CNT	B	D	CNT
John	Gold	4	54	100	4
Bill	Gold	8	54	80	2
			20	Silver	1






# Higher-Order IVM

## Theorem

For a variant of Relational Algebra with Aggregates, Higher-order IVM lowers the complexity of maintenance under single-tuple updates from complexity class AC0/TC0 to complexity class NC0.

C. Koch. Incremental query evaluation in a ring of databases. PODS 2010

## Practical system:



# TOASTER

## SQL QUERY COMPILER






*DBToaster: higher-order delta processing for dynamic, frequently fresh views.* C Koch et al. VLDB J. 23(2), 2014.



# TOASTER

# SQL QUERY COMPILER

## FRONT-END




Higher-order Incremental  
View Maintenance

Code Generation  
(C++, Scala, Spark)


# DBToaster: TPC-H BENCHMARK

## *Single-Tuple Incremental Stream Processing*




# DBToaster: TPC-H BENCHMARK

## *Single-Tuple Incremental Stream Processing*




# DBToaster: TPC-H BENCHMARK

## *Single-Tuple Incremental Stream Processing*



# DBToaster: TPC-H BENCHMARK

## *Single-Tuple Incremental Stream Processing*



# HIVM: disadvantage

$$Q = \pi_{AD}(R \bowtie S \bowtie T)$$

Maintain under all updates

Maintain base tables + query result ...

$R$

A	B	CNT
John	54	2
Mary	23	1
Bill	54	4

$S$

B	C	CNT
54	Gold	2
20	Silver	1

$T$

C	D	CNT
Gold	100	2
Gold	80	1
Silver	50	1
Bronze	20	4

$Q(D)$

A	D	CNT
John	100	8
John	80	4
Bill	100	16
Bill	80	8

... as well as query subresults as auxiliary views

$\pi_{AC}(R \bowtie S)$

A	C	CNT
John	Gold	4
Bill	Gold	8

$\pi_{BD}(S \bowtie T)$

B	D	CNT
54	100	4
54	80	2
20	Silver	1



Subresults can be of size  $|R| \times |S|$  resp.  $|S| \times |T|$

In general: can be bigger than  $|Q(D)|$

Not all subresults are useful to materialize

# IVM + HIVM: Disadvantage

$$Q = (R \bowtie S \bowtie T)$$

Full join query

(e.g., Complex Event Processing)

Maintain base tables + query result ...

*R*

A	B	CNT
John	54	2
Mary	23	1
Bill	54	4

*S*

B	C	CNT
54	Gold	2
20	Silver	1

*T*

C	D	CNT
Gold	100	2
Gold	80	1
Silver	50	1
Bronze	20	4

*Q(D)*

A	B	C	D	CNT
John	54	Gold	100	8
John	54	Gold	80	4
Bill	54	Gold	100	16
Bill	54	Gold	80	8


(... as well as query subresults as auxiliary views for HIVM)



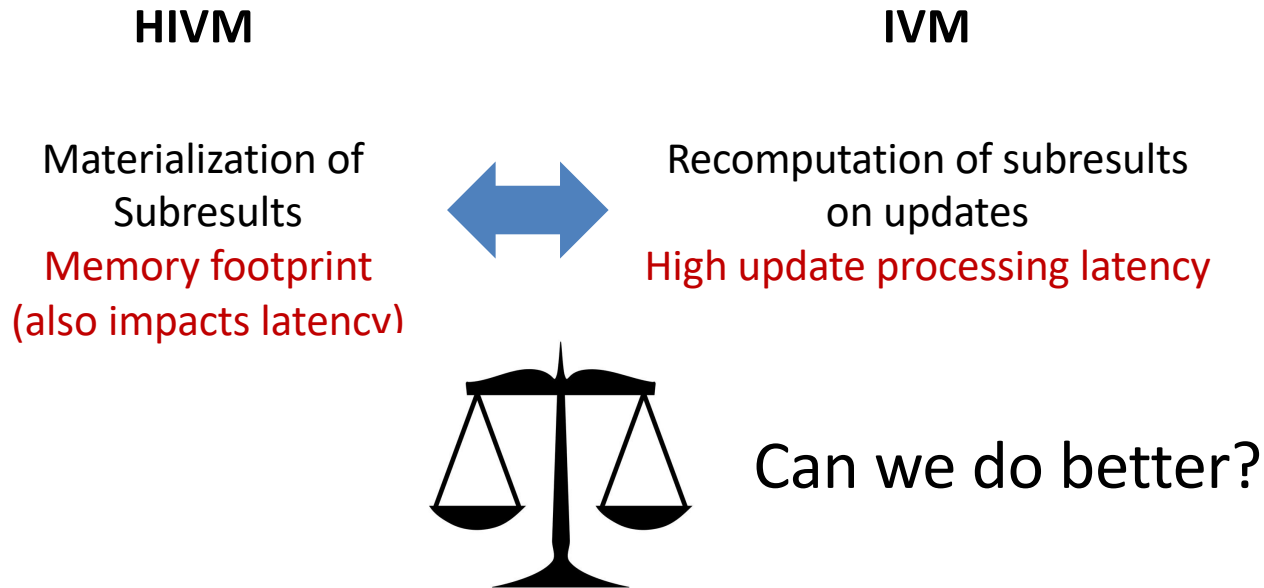
- Memory footprint:  $Q(D)$  can be  $|D|^2$
- Realistic to materialize in-memory ?



# Tradeoff between IVM and HIVM



# Tradeoff between IVM and HIVM



## Key Insight

Avoid storing the materialized views; instead store a compressed representation.

# Properties of Materialization of Q(D)

$$Q = (R \bowtie S \bowtie T)$$

A	B		B	C		C	D	
1	3		4	3		1	1	
3	4	⋈	3	5	⋈	3	4	=
2	4		6	5		4	5	
2	3		3	2		5	6	

A	B	C	D
3	4	3	4
2	4	3	4
1	3	5	6
2	3	5	6



Efficient  
enumeration  
of Q(DB)



Requires  $|Q(DB)|$  space,  
 $\Omega(N^2)$  in worst case

↓  
Array,  
Linked List,  
...

# Constant-delay enumeration

Enumerate set  $Q(D)$  with constant delay from a data structure  $M$ :

$Q(D)$

#	A	B	C	D
t1	a1	b1	c1	d1
t2	a2	b2	c2	d2
t3	a3	b3	c3	d3
.				
.				
.				
tn	an	bn	cn	dn

Constant time to enumerate first tuple

Constant time between tuples

Constant time between enumerating last tuple and end of enumeration process

Constant in **data complexity**: i.e., independent of  $|D|$  or  $|Q(D)|$ , but may depend on  $|Q|$ .

Comparable to enumerating from an in-memory array

≡ Streaming decompression algorithm

# Dynamic Yannakakis (DYN)

- Materialize a data structure that is:
  - **Succinct**: no larger than the database **D**
  - From which the query result can be **enumerated with constant delay**
  - Which can be **efficiently maintained** under updates

*The Dynamic Yannakakis Algorithm:  
Compact and Efficient Query Processing Under Updates.*  
M. Idris, M. Ugarte, S. Vansummeren. SIGMOD 2017

*Conjunctive Queries with Inequalities Under Updates.*  
M. Idris et al. PVLDB 11(7), 2018

# Dynamic Yannakakis (DYN)

- Materialize a data structure that is:
  - **Succinct**: no larger than the database **D**
  - From which the query result can be **enumerated with constant delay**
  - Which can be **efficiently maintained** under updates

To achieve this, DYN works only on **acyclic conjunctive queries** (but extends to deal with aggregation, negation).

## Conjunctive Queries (CQs)

Select-project-join queries with equi-joins only


$$Q = \pi_{ac} (R(a, b) \bowtie S(a, c) \bowtie T(c, d, e))$$

```
SELECT A, C, SUM(1)
FROM R, S, T
WHERE R.A = S.A and S.C = T.C
GROUP BY A, C
```

# Acyclic Queries

A CQ is acyclic if its body admits a Generalized Join Tree (GJT)\*

- GJTs for Acyclic CQs: A **node labelled tree**  $T$  where
  - Every **leaf is a relation** in the query



\*Does not depend on projections


$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$



# Acyclic Queries

A CQ is acyclic if its body admits a Generalized Join Tree (GJT)\*

- GJTs for Acyclic CQs: A **node labelled tree**  $T$  where
  - Every **leaf is a relation** in the query
  - Internal nodes are sets of attributes that are **subset of at least one** of its children (also called **guard**)




\*Does not depend on projections

$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

# Acyclic Queries

A CQ is acyclic if its body admits a Generalized Join Tree (GJT)\*

- GJTs for Acyclic CQs: A **node labelled tree**  $T$  where
  - Every **leaf is a relation** in the query
  - Internal nodes are sets of attributes that are **subset of at least one** of its children (also called **guard**)




\*Does not depend on projections

$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

# Acyclic Queries

A CQ is acyclic if its body admits a Generalized Join Tree (GJT)\*

- GJTs for Acyclic CQs: A **node labelled tree**  $T$  where
  - Every **leaf is a relation** in the query
  - Internal nodes are sets of attributes that are **subset of at least one** of its children (also called **guard**)
  - Every variable in the tree induces a **connected subtree** in  $T$




\*Does not depend on  
projections

$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

# Acyclic Queries

A CQ is acyclic if its body admits a Generalized Join Tree (GJT)\*

- GJTs for Acyclic CQs: A **node labelled tree**  $T$  where
  - Every **leaf is a relation** in the query
  - Internal nodes are sets of attributes that are **subset of at least one** of its children (also called **guard**)
  - Every variable in the tree induces a **connected subtree** in  $T$





\*Does not depend on  
projections

$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

## A cyclic query: the triangle query

$$Q = R(a, b) \bowtie S(b, c) \bowtie T(a, c)$$



# Objectives of Dyanmic Yannakakis (DYN)

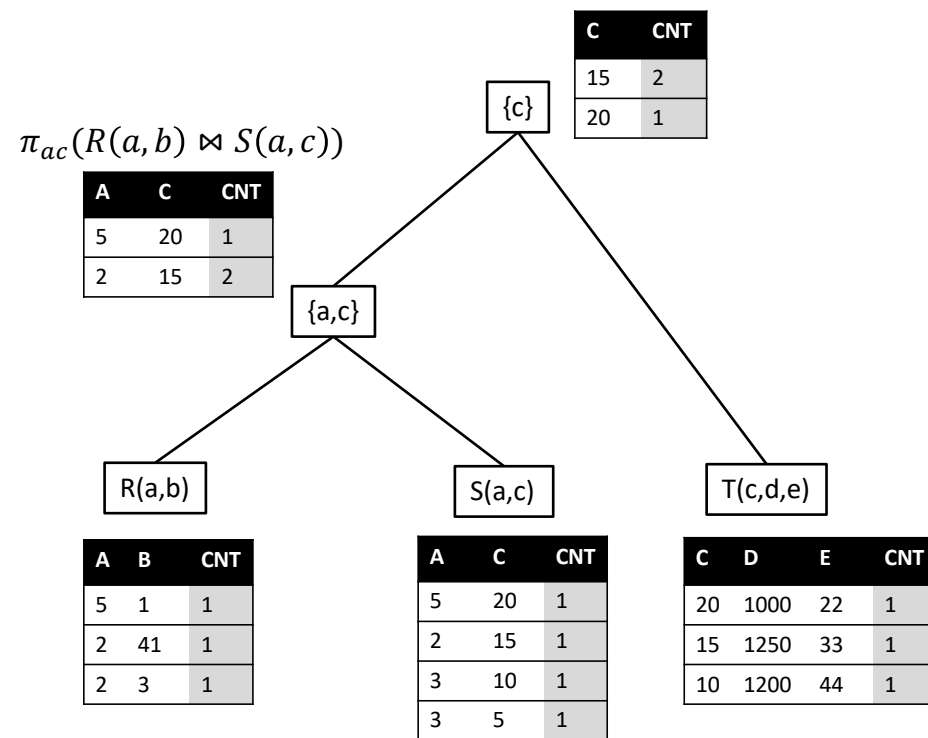
- Materialize a data structure that is:
  - **Succinct**: no larger than the database D
  - From which the query result can be **enumerated with constant delay**
  - Which can be **efficiently maintained** under updates

## T-reduct: Compressed representation based on GJT

$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

$$\pi_c(R(a, b) \bowtie S(a, c) \bowtie T(c, d, e))$$

- Bottom-up semi-Join reduction
- Linear in the size of database



# Objectives of Dyanmic Yannakakis (DYN)

- Materialize a data structure that is:
  - **Succinct**: no larger than the database **D**
  - From which the query result can be **enumerated with constant delay**
  - Which can be **efficiently maintained** under updates



# T-reduct: Enumeration

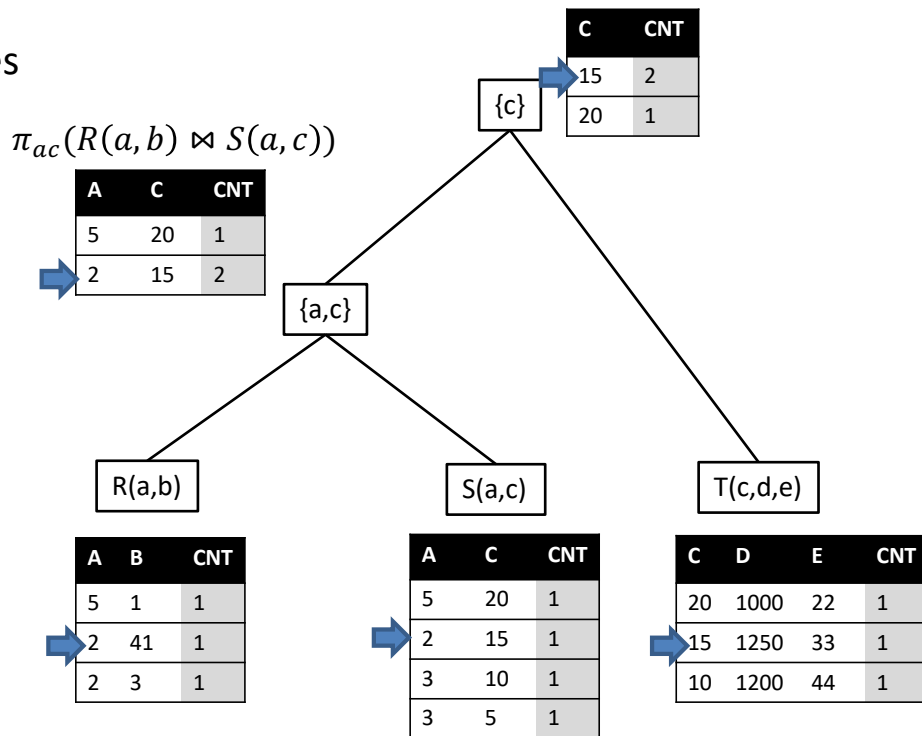
$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

$$\pi_c(R(a, b) \bowtie S(a, c) \bowtie T(c, d, e))$$

- Top-down: find “compatible” tuples

Output:

A	B	C	D	E	CNT
2	41	15	1250	33	1



# T-reduct: Enumeration


$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

$$\pi_c(R(a, b) \bowtie S(a, c) \bowtie T(c, d, e))$$

- Top-down: find “compatible” tuples

Output:

A	B	C	D	E	CNT
2	41	15	1250	33	1
2	3	15	1250	33	1



# T-reduct: Enumeration

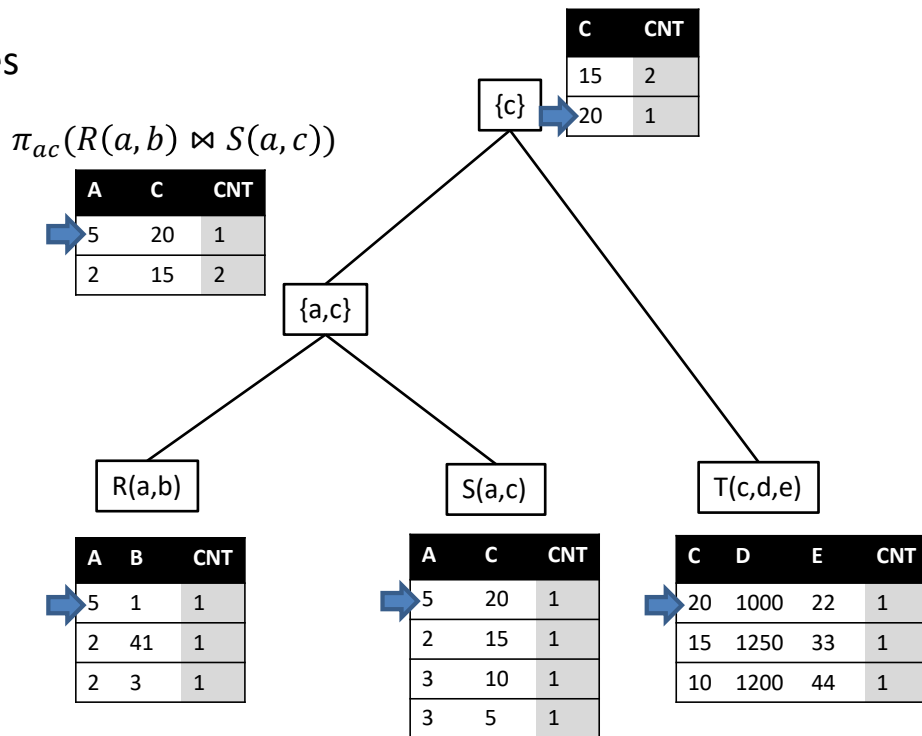
$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

$$\pi_c(R(a, b) \bowtie S(a, c) \bowtie T(c, d, e))$$

- Top-down: find “compatible” tuples

Output:

A	B	C	D	E	CNT
2	41	15	1250	33	1
2	3	15	1250	33	1
5	1	20	1000	22	1



# T-reduct: Enumeration

$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

$$\pi_{cd}(R(a, b) \bowtie S(a, c) \bowtie T(c, d, e))$$

- Top-down: find “compatible” tuples
- Add indexes to make this efficient

$\pi_{cd}(R(a, b) \bowtie S(a, c))$

C	CNT
15	2
20	1



- Assuming that index access is  $O(1)$ , this is enumeration with constant delay

Output:

A	B	C	D	E	CNT
2	41	15	1250	33	1
2	3	15	1250	33	1
5	1	20	1000	22	1

$R(a, b)$

A	B	CNT
5	1	1
2	41	1
	3	1

$S(a, c)$

A	C	CNT
5	20	1
2	15	1
3	10	1
3	5	1


$T(c, d, e)$

C	D	E	CNT
15	1250	33	1
20	1000	22	1
10	1200	44	1

# What about projections ?

$$Q = \pi_{a,c}(R(a,b) \bowtie S(a,c) \bowtie T(c,d,e))$$

Compatible tree



We can still enumerate  
with constant delay!

# Objectives of Dyanmic Yannakakis (DYN)

- Materialize a data structure that is:
  - **Succinct**: no larger than the database D
  - From which the query result can be **enumerated with constant delay**
  - Which can be **efficiently maintained** under updates

# T-reduct: Update processing

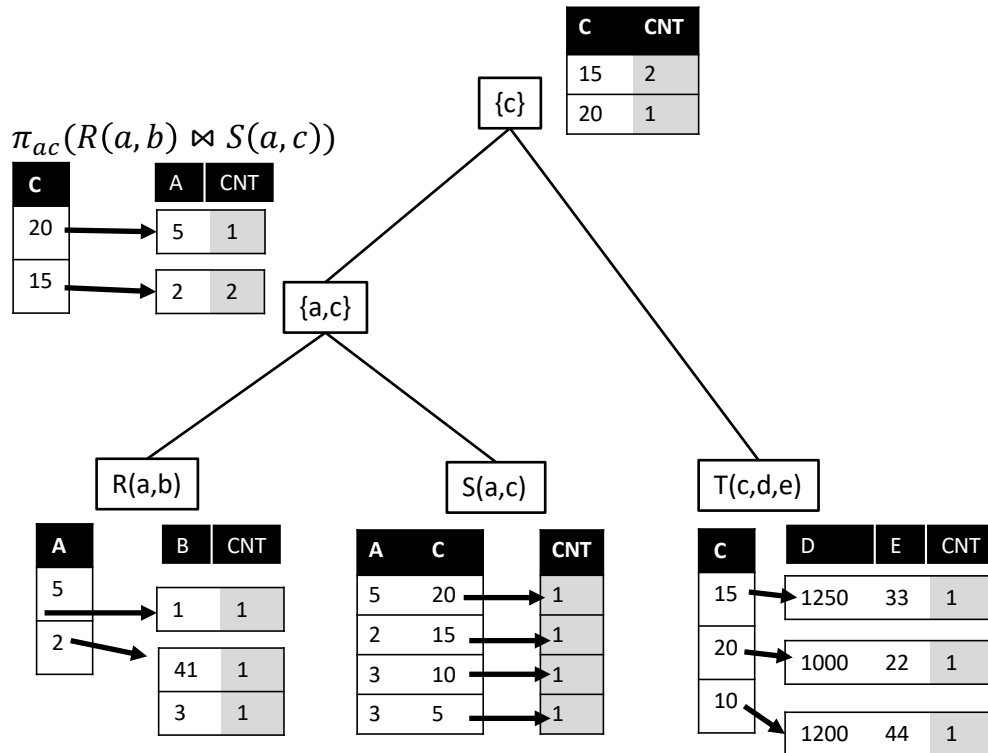
$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

$$\pi_c(R(a, b) \bowtie S(a, c) \bowtie T(c, d, e))$$

- Bottom-up: propagate update through semi-join reduction (view each node as IVM)

$\Delta S$

A	C	CNT
5	10	1
5	20	-1



# T-reduct: Update processing

$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

$$\pi_c(R(a, b) \bowtie S(a, c) \bowtie T(c, d, e))$$

- Bottom-up: propagate update through semi-join reduction (view each node as IVM)

$$\Delta_{ac} = \pi_{ac}(R(a, b) \bowtie \Delta S(a, c))$$

A	C	CNT
5	10	1
5	20	-1

$$\Delta S$$

A	C	CNT
5	10	1
5	20	-1

$$\pi_{ac}(R(a, b) \bowtie S(a, c))$$

C	A	CNT
20	5	1
15	2	2

{a,c}

R(a,b)

A	B	CNT
5	1	1
2	41	1
	3	1

S(a,c)

A	C	CNT
5	20	0
2	15	1
3	10	1
3	5	1
5	10	1

T(c,d,e)

C	D	E	CNT
15	1250	33	1
20	1000	22	1
10	1200	44	1

C	CNT
15	2
20	1



# T-reduct: Update processing

$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

- Bottom-up: propagate update through semi-join reduction (view each node as IVM)

$$\Delta_{ac} = \pi_{ac}(R(a, b) \bowtie \Delta S(a, c))$$

A	C	CNT
5	10	1
5	20	-1

$\Delta S$

A	C	CNT
5	10	1
5	20	-1

$$\pi_{ac}(R(a, b) \bowtie S(a, c))$$

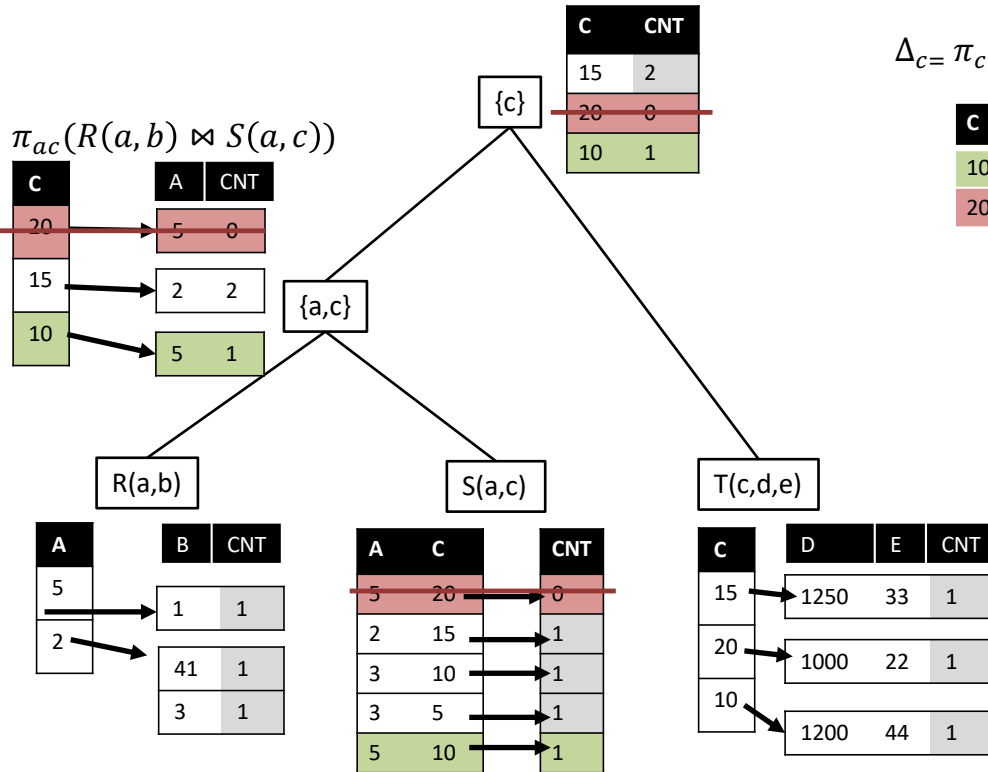
C	A	CNT
20	5	0
15	2	2
10	5	1

$$\pi_c(R(a, b) \bowtie S(a, c) \bowtie T(c, d, e))$$

C	CNT
15	2
20	0
10	1

$$\Delta_c = \pi_c(\Delta_{ac} \bowtie T(c, d, e))$$

C	CNT
10	1
20	-1



# T-reduct: Update processing

$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

- Index guards on attributes shared with sibling

$$\Delta_{ac} = \pi_{ac}(\Delta R(a, b) \bowtie S(a, c))$$

A	C	CNT
3	10	1
3	5	1

$\Delta R$

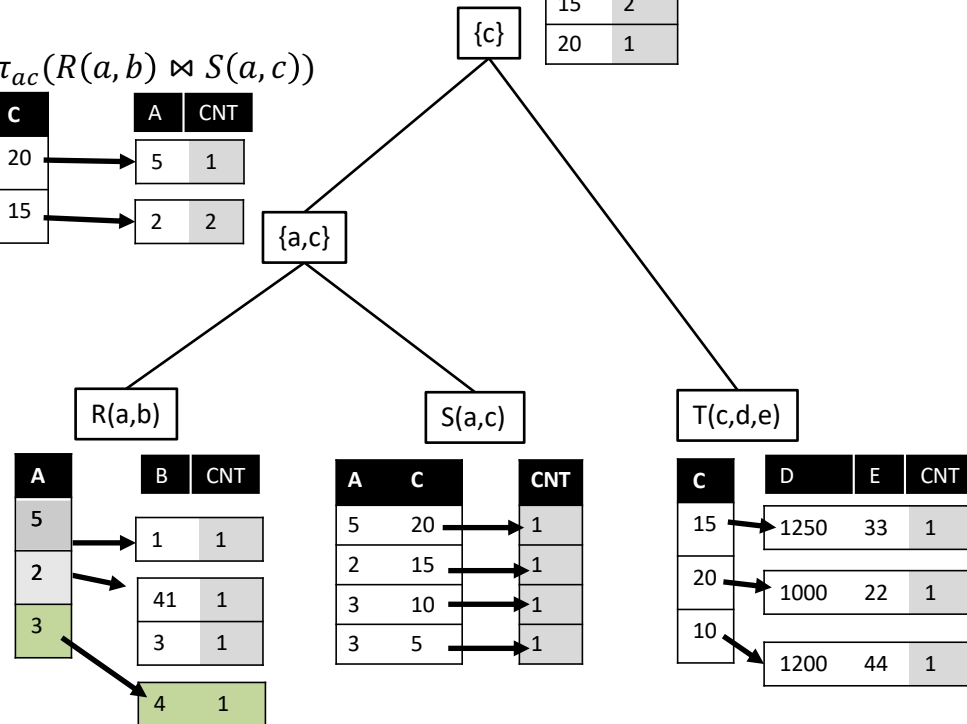
A	B	CNT
3	4	1

$$\pi_{ac}(R(a, b) \bowtie S(a, c))$$

C	A	CNT
20	5	1
15	2	2

$$\pi_c(R(a, b) \bowtie S(a, c) \bowtie T(c, d, e))$$

C	CNT
15	2
20	1



# T-reduct: Update processing

$$Q = R(a, b) \bowtie S(a, c) \bowtie T(c, d, e)$$

- Index guards on attributes shared with sibling

$$\pi_c(R(a, b) \bowtie S(a, c) \bowtie T(c, d, e))$$

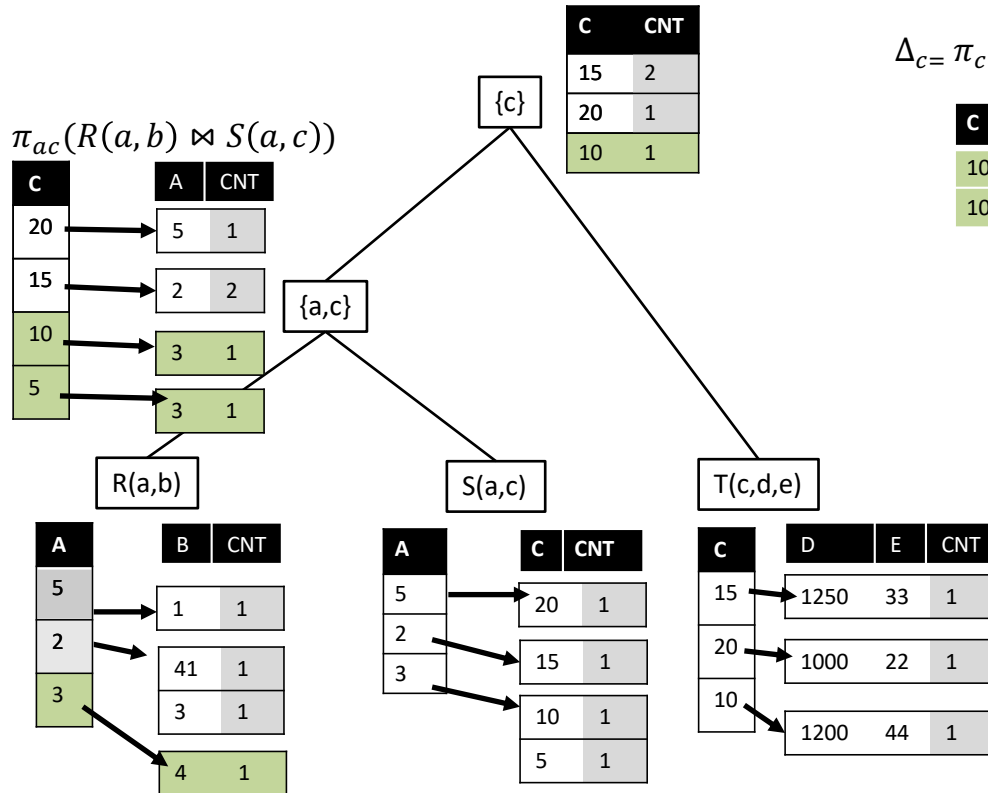
$$\Delta_c = \pi_c(\Delta_{ac} \bowtie T(c, d, e))$$

$$\Delta_{ac} = \pi_{ac}(\Delta R(a, b) \bowtie S(a, c))$$

A	C	CNT
3	10	1
3	5	1


$\Delta R$

A	B	CNT
3	4	1




# DYN in practice

100% = DBToaster




- Up to two order of magnitude faster
- Consumes up to one order of magnitude less memory

# Enumeration

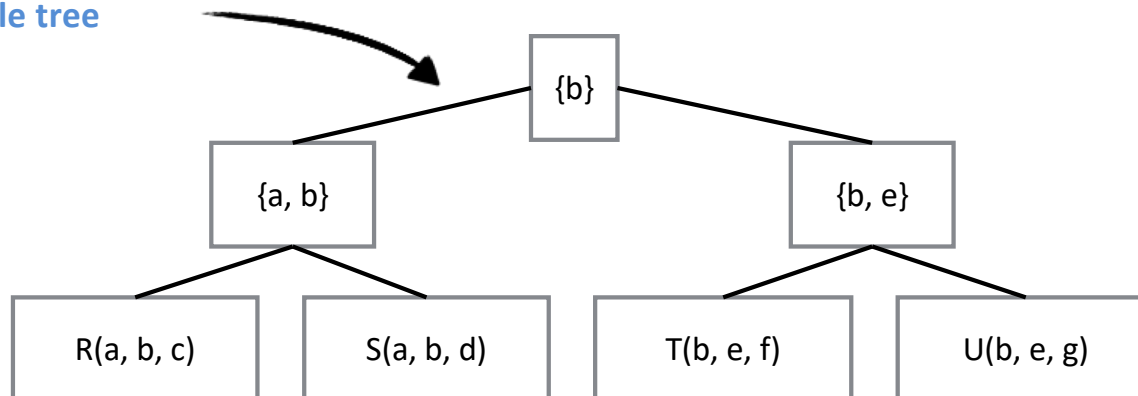


# A note about complexity

Assume  $|\Delta D|$  is constant, is the required update time also constant?



Simple tree



# A note about complexity



- Ideal IVM algorithm allows, for any query  $Q$ 
  - Constant delay enumeration of  $Q(D)$
  - Constant-time update processing if  $|\Delta D|$  is constant

[Berkholz et al., PODS 2017; ICDT 2018]

Conjunctive query  $Q$  supports constant-delay enumeration  
after constant-time updates




$Q$  is *q-hierarchical*

(under certain complexity-theoretical assumptions)

# Q-hierarchical CQs

A CQ is **q-hierarchical** if its body has a **q-tree**\*

$$\pi_{a,b,e}(R(a,b,c) \bowtie S(a,b,d) \bowtie T(b,e,f) \bowtie U(b,e,g))$$



Atom condition: each atom must induce a directed path, starting from the root

Projection condition: the projection attributes must induce a connected subtree that contains the root

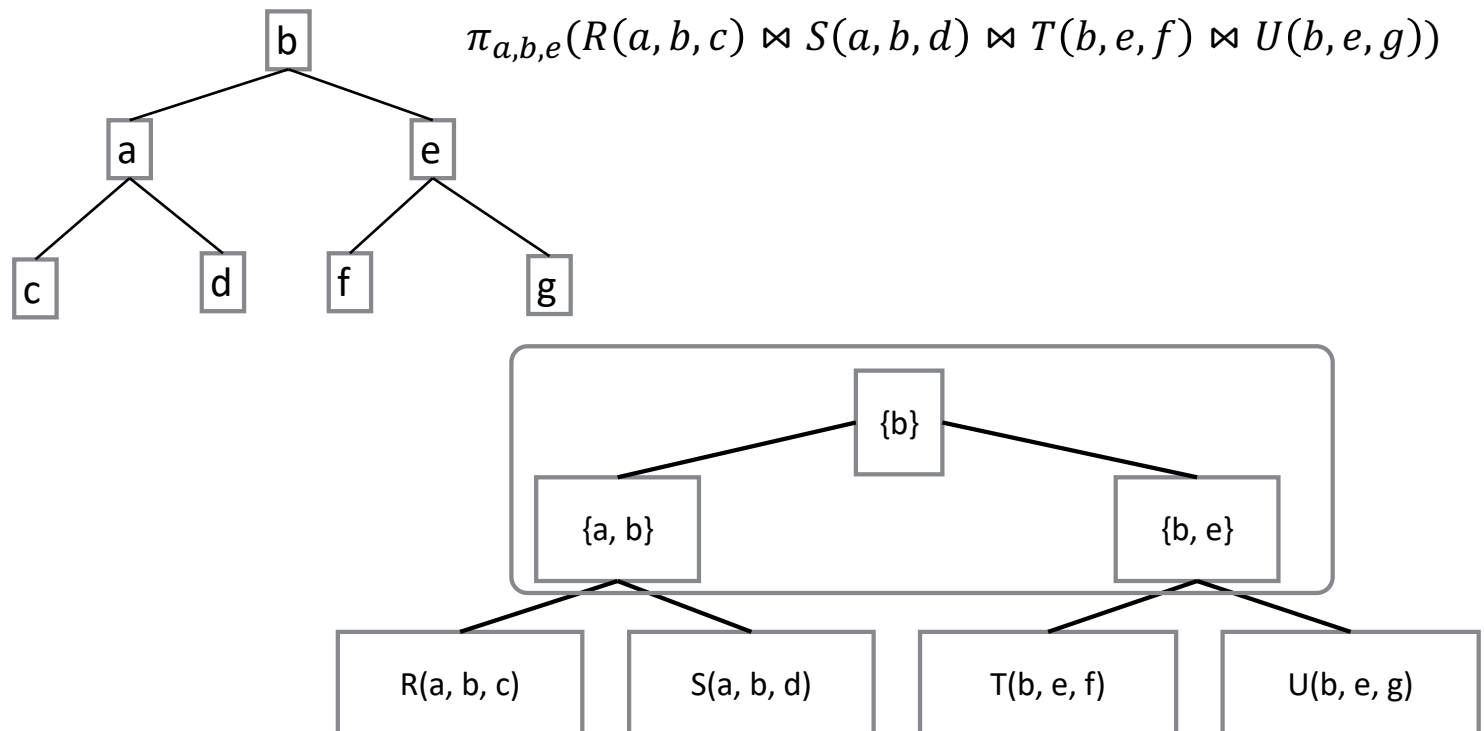
\*Does depend on projections



# Q-hierarchical CQs

## Theorem

A CQ is q-hierarchical if and only if it has a generalized join tree that is both simple and compatible with the projection



# Q-hierarchical CQs

## Theorem

A CQ is q-hierarchical if and only if it has a generalized join tree that is both simple and compatible with the projection

## Corollary

DYN provides constant-delay enumeration after constant-time updates precisely for the class of q-hierarchical queries.

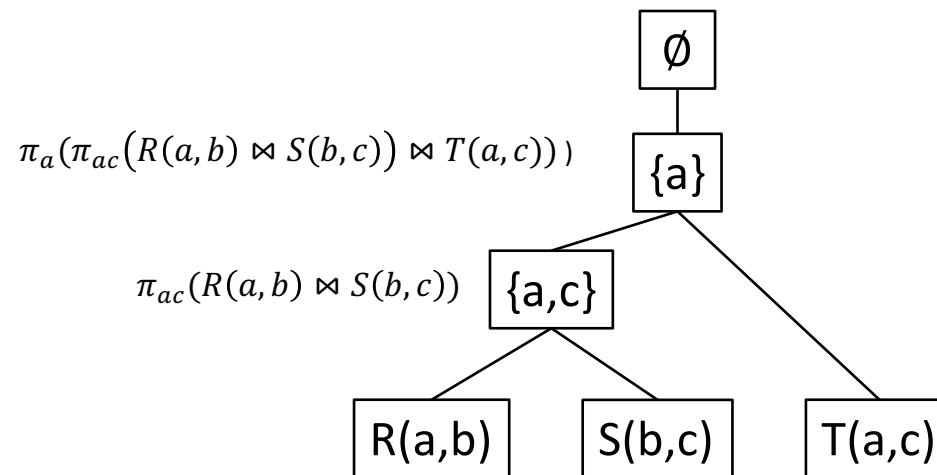
Matches the theoretical lower bound.



# What about Cyclic Queries?

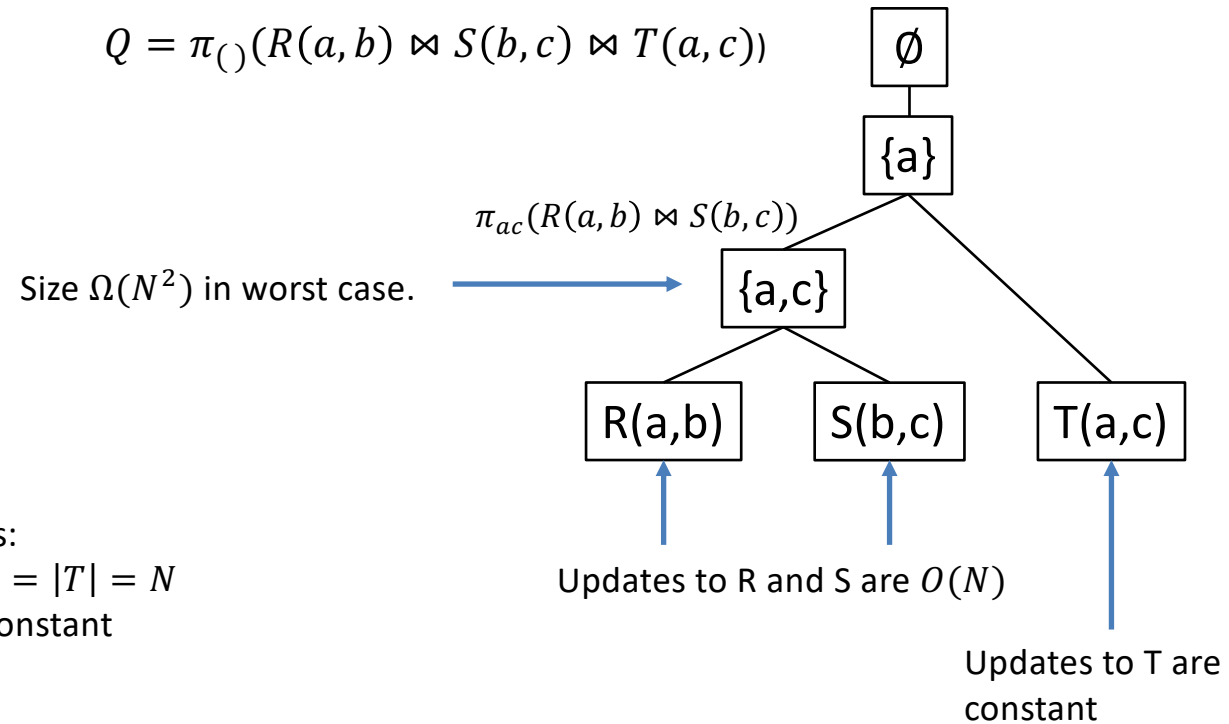
- Use **View Trees**, which relax constraints on GJTs:
  - interior nodes need not have guards
  - Give up connectedness condition

$$Q = \pi_{()}(R(a, b) \bowtie S(b, c) \bowtie T(a, c))$$



# What about Cyclic Queries?

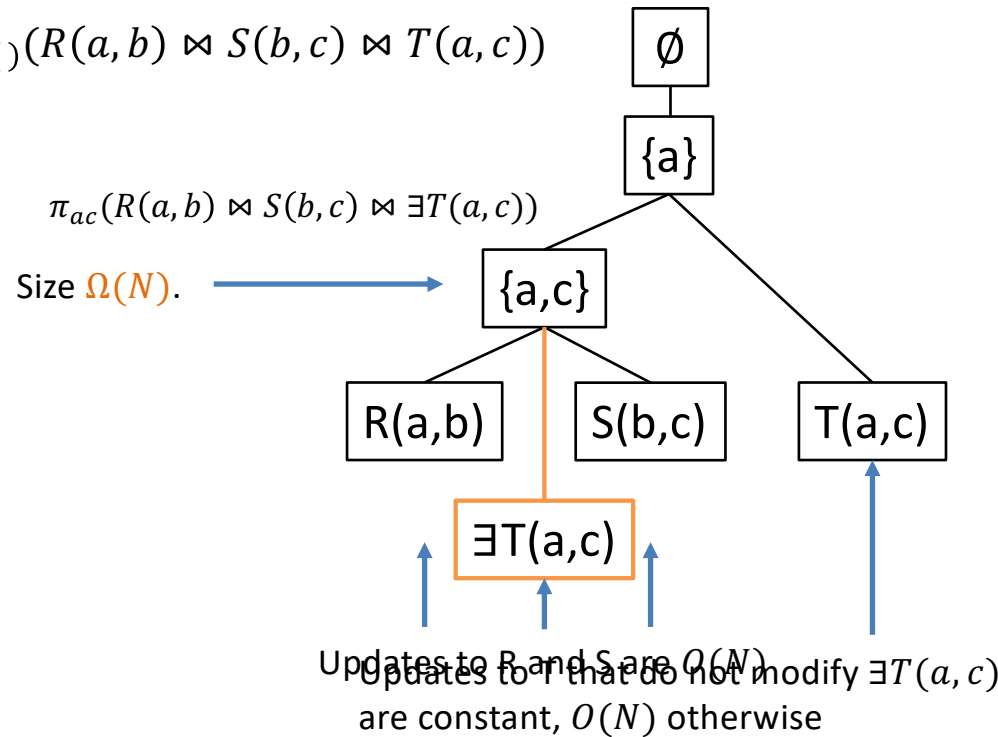
- Use View Trees, relaxing constraints on GJTs
- Nodes without guard may be superlinear in  $|D|$  but may help processing to some updates



# What about Cyclic Queries?

- Use View Trees, relaxing constraints on GJT's
- Nodes without guard may be superlinear in  $|D|$  but may help processing to some updates
- Add indicator projections to minimize space

$$Q = \pi_{()}(R(a,b) \bowtie S(b,c) \bowtie T(a,c))$$



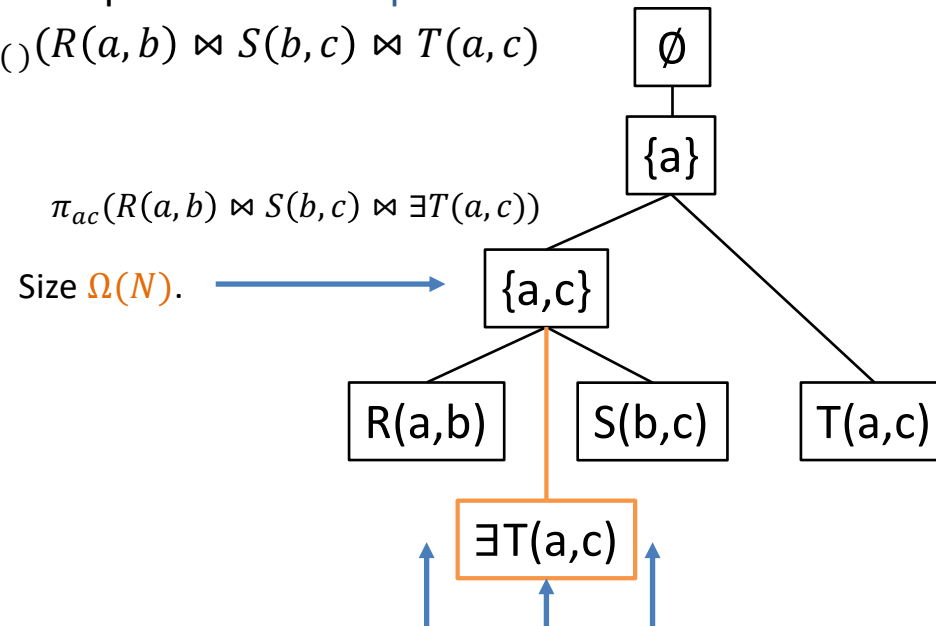
Assumptions:

- $|R| = |S| = |T| = N$
- $|\Delta D| = \text{constant}$

# What about Cyclic Queries?

- Use View Trees, relaxing constraints on GJT's
- Nodes without guard may be superlinear in  $|D|$  but may help processing to some updates
- Add indicator projections to minimize space and **bulk update time**

$$Q = \pi_{\langle \rangle}(R(a, b) \bowtie S(b, c) \bowtie T(a, c))$$



Assumptions:

- $|R| = |S| = |T| = N$
- Bulk updates,  $|\Delta D| = N$

Bulk updates to  $R$ ,  $S$ , or  $\exists T$  are  $\Theta(N^{3/2})$  by use of worst-case optimal join algorithms

# What about Cyclic Queries

- This approach proposed in:

*Incremental View Maintenance with Triple Lock Factorization Benefits.*

Milos Nikolic, Dan Olteanu:

SIGMOD Conference 2018: 365-380

- F-IVM features:
  - View trees instead of GJTs for HIVM-based processing also allowing cyclic queries
  - Processing of complex aggregations (see later)
  - Exploiting factorized representations of updates and results (see later)

# Main Algorithmic Ideas

1. IVM  $\equiv$  processing of delta queries

2. Materialize results of subqueries in addition to the actual query result

➔ 3. Exploit data skew

Time

1993

1997

2009

2018




# Exploiting Data Skew

- The maintenance approaches considered so far exploit query structure but not data skew.
- These approaches do not achieve worst-case optimal update and answer times in general.
  - Exception: q-hierarchical queries
- We present a maintenance approach that takes data skew into account and admits:
  - Worst-case optimal update and answer times
  - Time-space trade-off

Counting Triangles under Updates in Worst-Case Optimal Time.  
Ahmet Kara, Hung Q. Ngo, Milos Nikolic, Dan Olteanu, Haozhe Zhang.  
To appear in ICDT 2019

# Example: The Triangle Count

Maintain the triangle count  $Q$   
under single-tuple updates to  $R$ ,  $S$ , and  $T$ !



$Q$  counts the number of tuples  
in the join of  $R$ ,  $S$ , and  $T$ .

$$Q = \pi_{()} [R(a, b) \bowtie S(b, c) \bowtie T(c, a)]$$

# Updates to the Triangle Count

$R$			$S$			$T$		
$A$	$B$		$B$	$C$		$C$	$A$	
$a_1$	$b_1$	2	$b_1$	$c_1$	2	$c_1$	$a_1$	1
$a_2$	$b_1$	3	$b_1$	$c_2$	1	$c_2$	$a_1$	3
						$c_2$	$a_2$	3

# Updates to the Triangle Count

$R$	$S$	$T$	$R \bowtie S \bowtie T$
$A \ B \mid$	$B \ C \mid$	$C \ A \mid$	$A \ B \ C \mid$
$a_1 \ b_1 \mid 2$	$b_1 \ c_1 \mid 2$	$c_1 \ a_1 \mid 1$	$a_1 \ b_1 \ c_1 \mid 2 \cdot 2 \cdot 1 = 4$
$a_2 \ b_1 \mid 3$	$b_1 \ c_2 \mid 1$	$c_2 \ a_1 \mid 3$	
		$c_2 \ a_2 \mid 3$	

# Updates to the Triangle Count

$R$	$S$	$T$	$R \bowtie S \bowtie T$
$A \mid B \mid$	$B \mid C \mid$	$C \mid A \mid$	$A \mid B \mid C \mid$
$a_1 \mid b_1 \mid 2$	$b_1 \mid c_1 \mid 2$	$c_1 \mid a_1 \mid 1$	$a_1 \mid b_1 \mid c_1 \mid 2 \cdot 2 \cdot 1 = 4$
$a_2 \mid b_1 \mid 3$	$b_1 \mid c_2 \mid 1$	$c_2 \mid a_1 \mid 3$	$a_1 \mid b_1 \mid c_2 \mid 2 \cdot 1 \cdot 3 = 6$
		$c_2 \mid a_2 \mid 3$	$a_2 \mid b_1 \mid c_2 \mid 3 \cdot 1 \cdot 3 = 9$

# Updates to the Triangle Count

$R$	$S$	$T$	$R \bowtie S \bowtie T$
$A \mid B \mid$	$B \mid C \mid$	$C \mid A \mid$	$A \mid B \mid C \mid$
$a_1 \mid b_1 \mid 2$	$b_1 \mid c_1 \mid 2$	$c_1 \mid a_1 \mid 1$	$a_1 \mid b_1 \mid c_1 \mid 2 \cdot 2 \cdot 1 = 4$
$a_2 \mid b_1 \mid 3$	$b_1 \mid c_2 \mid 1$	$c_2 \mid a_1 \mid 3$	$a_1 \mid b_1 \mid c_2 \mid 2 \cdot 1 \cdot 3 = 6$
		$c_2 \mid a_2 \mid 3$	$a_2 \mid b_1 \mid c_2 \mid 3 \cdot 1 \cdot 3 = 9$



$Q(\mathbf{D})$
$\emptyset \mid$
$( ) \mid 4 + 6 + 9 = 19$

# Updates to the Triangle Count

$R$	$S$	$T$	$R \bowtie S \bowtie T$
$A \ B \mid$	$B \ C \mid$	$C \ A \mid$	$A \ B \ C \mid$
$a_1 \ b_1 \mid 2$	$b_1 \ c_1 \mid 2$	$c_1 \ a_1 \mid 1$	$a_1 \ b_1 \ c_1 \mid 2 \cdot 2 \cdot 1 = 4$
$a_2 \ b_1 \mid 3$	$b_1 \ c_2 \mid 1$	$c_2 \ a_1 \mid 3$	$a_1 \ b_1 \ c_2 \mid 2 \cdot 1 \cdot 3 = 6$
		$c_2 \ a_2 \mid 3$	$a_2 \ b_1 \ c_2 \mid 3 \cdot 1 \cdot 3 = 9$



$\Delta R(a, b)$
$A \ B \mid$
$a_2 \ b_1 \mid -2$



$Q(\mathbf{D})$
$\emptyset \mid$
$() \mid 4 + 6 + 9 = 19$

# Updates to the Triangle Count

$R$	$S$	$T$	$R \bowtie S \bowtie T$
$A \ B \mid$	$B \ C \mid$	$C \ A \mid$	$A \ B \ C \mid$
$a_1 \ b_1 \mid 2$	$b_1 \ c_1 \mid 2$	$c_1 \ a_1 \mid 1$	$a_1 \ b_1 \ c_1 \mid 2 \cdot 2 \cdot 1 = 4$
$a_2 \ b_1 \mid 3$	$b_1 \ c_2 \mid 1$	$c_2 \ a_1 \mid 3$	$a_1 \ b_1 \ c_2 \mid 2 \cdot 1 \cdot 3 = 6$
		$c_2 \ a_2 \mid 3$	$a_2 \ b_1 \ c_2 \mid 3 \cdot 1 \cdot 3 = 9$



$\Delta R(a, b)$
$A \ B \mid$
$a_2 \ b_1 \mid -2$



$Q(D)$
$\emptyset \mid$
$() \mid 4 + 6 + 9 = 19$



# Updates to the Triangle Count

$R$	$S$	$T$	$R \bowtie S \bowtie T$
$A \ B \mid$	$B \ C \mid$	$C \ A \mid$	$A \ B \ C \mid$
$a_1 \ b_1 \mid 2$	$b_1 \ c_1 \mid 2$	$c_1 \ a_1 \mid 1$	$a_1 \ b_1 \ c_1 \mid 2 \cdot 2 \cdot 1 = 4$
$a_2 \ b_1 \mid 1$	$b_1 \ c_2 \mid 1$	$c_2 \ a_1 \mid 3$	$a_1 \ b_1 \ c_2 \mid 2 \cdot 1 \cdot 3 = 6$
		$c_2 \ a_2 \mid 3$	$a_2 \ b_1 \ c_2 \mid 3 \cdot 1 \cdot 3 = 9$



$\Delta R(a, b)$
$A \ B \mid$
$a_2 \ b_1 \mid -2$



$Q(D)$
$\emptyset \mid$
$() \mid 4 + 6 + 9 = 19$

# Updates to the Triangle Count

$R$	$S$	$T$	$R \bowtie S \bowtie T$
$A \ B \mid$	$B \ C \mid$	$C \ A \mid$	$A \ B \ C \mid$
$a_1 \ b_1 \mid 2$	$b_1 \ c_1 \mid 2$	$c_1 \ a_1 \mid 1$	$a_1 \ b_2 \ c_2 \mid 2 \cdot 2 \cdot 1 = 4$
$a_2 \ b_1 \mid 1$	$b_1 \ c_2 \mid 1$	$c_2 \ a_1 \mid 3$	$a_1 \ b_1 \ c_2 \mid 2 \cdot 1 \cdot 3 = 6$
		$c_2 \ a_2 \mid 3$	$a_2 \ b_1 \ c_3 \mid 3 \cdot 1 \cdot 3 = 9$



$\Delta R(a, b)$
$A \ B \mid$
$a_2 \ b_1 \mid -2$



$Q(D)$
$\emptyset \mid$
$() \mid 4 + 6 + 9 = 19$

# Updates to the Triangle Count

$R$	$S$	$T$	$R \bowtie S \bowtie T$
$A \ B \mid$	$B \ C \mid$	$C \ A \mid$	$A \ B \ C \mid$
$a_1 \ b_1 \mid 2$	$b_1 \ c_1 \mid 2$	$c_1 \ a_1 \mid 1$	$a_1 \ b_2 \ c_2 \mid 2 \cdot 2 \cdot 1 = 4$
$a_2 \ b_1 \mid 1$	$b_1 \ c_2 \mid 1$	$c_2 \ a_1 \mid 3$	$a_1 \ b_1 \ c_2 \mid 2 \cdot 1 \cdot 3 = 6$
		$c_2 \ a_2 \mid 3$	$a_2 \ b_1 \ c_3 \mid 1 \cdot 1 \cdot 3 = 3$



$\Delta R(a, b)$
$A \ B \mid$
$a_2 \ b_1 \mid -2$



$Q(D)$
$\emptyset \mid$
$() \mid 4 + 6 + 9 = 19$

# Updates to the Triangle Count

$R$	$S$	$T$	$R \bowtie S \bowtie T$
$A \ B \mid$	$B \ C \mid$	$C \ A \mid$	$A \ B \ C \mid$
$a_1 \ b_1 \mid 2$	$b_1 \ c_1 \mid 2$	$c_1 \ a_1 \mid 1$	$a_1 \ b_2 \ c_2 \mid 2 \cdot 2 \cdot 1 = 4$
$a_2 \ b_1 \mid 1$	$b_1 \ c_2 \mid 1$	$c_2 \ a_1 \mid 3$	$a_1 \ b_1 \ c_2 \mid 2 \cdot 1 \cdot 3 = 6$
		$c_2 \ a_2 \mid 3$	$a_2 \ b_1 \ c_3 \mid 1 \cdot 1 \cdot 3 = 3$

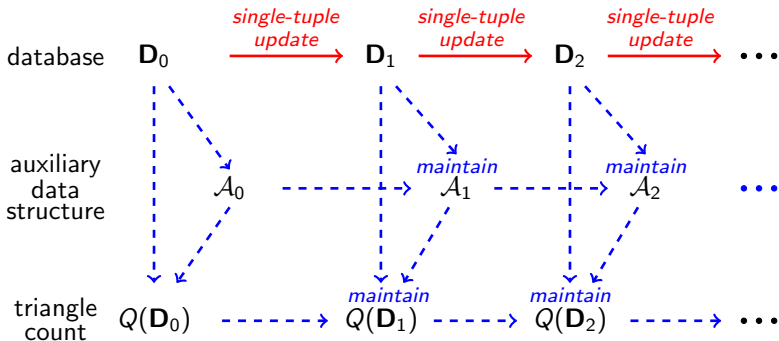


$\Delta R(a, b)$
$A \ B \mid$
$a_2 \ b_1 \mid -2$



$Q(D)$
$\emptyset \mid$
$( ) \mid 4 + 6 + 3 = 13$

# The Considered Maintenance Problem



Given a current database  $\mathbf{D}$  and a single-tuple update, what are the time and space complexities for maintaining  $Q(\mathbf{D})$ ?

# Much Ado about Triangles

The Triangle Query Served as Milestone in Many Fields

- Worst-case optimal join algorithms [*Algorithmica* 1997, *SIGMOD R.* 2013]
- Parallel query evaluation [*Found. & Trends DB* 2018]
- Randomized approximation in static settings [*FOCS* 2015]
- Randomized approximation in data streams  
[*SODA* 2002, *COCOON* 2005, *PODS* 2006, *PODS* 2016, *Theor. Comput. Sci.* 2017]

Intensive Investigation of Answering Queries under Updates

- Theoretical developments [*PODS* 2017, *ICDT* 2018]
- Systems developments [*F. & T. DB* 2012, *VLDB J.* 2014, *SIGMOD* 2017, 2018]
- Lower bounds [*STOC* 2015, *ICM* 2018]

So far: **No** dynamic algorithm maintaining the  
**exact triangle count** in **worst-case optimal** time!

# Naïve Maintenance

*"Compute from scratch!"*

$$\begin{aligned} \pi_{()} \left[ \underbrace{(R(a, b) + \Delta R(a, b))}_{\text{newR}} \bowtie S(b, c) \bowtie T(c, a) \right] \\ = \\ \pi_{()} \left[ \text{newR}(a, b) \bowtie S(b, c) \bowtie T(c, a) \right] \end{aligned}$$

## Maintenance Complexity

- Time:  $\mathcal{O}(|\mathbf{D}|^{1.5})$  using worst-case optimal join algorithms
- Space:  $\mathcal{O}(|\mathbf{D}|)$  to store input relations

# Classical IVM

*“Compute the difference!”*

Let  $\Delta R(a, b) = \{(a', b') \mapsto m\}$

$$\begin{aligned} \pi_{()} \left[ \left( R(a, b) + \Delta R(a, b) \right) \bowtie S(b, c) \bowtie T(c, a) \right] \\ = \\ \pi_{()} \left[ R(a, b) \bowtie S(b, c) \bowtie T(c, a) \right] \\ + \\ \pi_{()} \left[ \Delta R(a, b) \bowtie \sigma_{b=b'} S(b, c) \bowtie \sigma_{a=a'} T(c, a) \right] \end{aligned}$$

Maintenance Complexity

- Time:  $\mathcal{O}(|\mathbf{D}|)$  to intersect  $C$ -values from  $S$  and  $T$
- Space:  $\mathcal{O}(|\mathbf{D}|)$  to store input relations



# Higher Order IVM

*“Compute the difference by using pre-materialized views!”*

Let  $\Delta R(a, b) = \{(a', b') \mapsto m\}$

Pre-materialize  $V_{ST}(b, a) = \pi_{b,a} S(b, c) \bowtie T(c, a)!$

$$\begin{aligned} \pi_{()} \left[ (R(a, b) + \Delta R(a, b)) \bowtie S(b, c) \bowtie T(c, a) \right] \\ = \\ \pi_{()} \left[ R(a, b) \bowtie S(b, c) \bowtie T(c, a) \right] \\ + \\ \pi_{()} \left[ \Delta R(a, b) \bowtie \sigma_{a=a', b=b'} V_{ST}(b, a) \right] \end{aligned}$$

Maintenance Complexity

- Time for updates to  $R$ :  $\mathcal{O}(1)$  to look up in  $V_{ST}$
- Time for updates to  $S$  and  $T$ :  $\mathcal{O}(|\mathbf{D}|)$  to maintain  $V_{ST}$
- Space:  $\mathcal{O}(|\mathbf{D}|^2)$  to store input relations and  $V_{ST}$  (improvable to  $\mathcal{O}(|\mathbf{D}|^{1.5})$ )

# Closing the Complexity Gap

Complexity bounds for the maintenance of the triangle count

## Known Upper Bound

Maintenance Time:  $\mathcal{O}(|\mathbf{D}|)$

Space:  $\mathcal{O}(|\mathbf{D}|)$

## Known Lower Bound

Amortized maintenance time: **not**  $\mathcal{O}(|\mathbf{D}|^{0.5-\gamma})$  for any  $\gamma > 0$   
(under reasonable complexity theoretic assumptions)

# Closing the Complexity Gap

Complexity bounds for the maintenance of the triangle count

## Known Upper Bound

Maintenance Time:  $\mathcal{O}(|\mathbf{D}|)$

Space:  $\mathcal{O}(|\mathbf{D}|)$

Can the triangle count  
be maintained in  
sublinear time?

## Known Lower Bound

Amortized maintenance time: **not**  $\mathcal{O}(|\mathbf{D}|^{0.5-\gamma})$  for any  $\gamma > 0$   
(under reasonable complexity theoretic assumptions)

# Closing the Complexity Gap

Complexity bounds for the maintenance of the triangle count

## Known Upper Bound

Maintenance Time:  $\mathcal{O}(|\mathbf{D}|)$

Space:  $\mathcal{O}(|\mathbf{D}|)$

**Yes!**

Can the triangle count  
be maintained in  
sublinear time?

We propose:  $\text{IVM}^\epsilon$

Amortized maintenance time:

$\mathcal{O}(|\mathbf{D}|^{0.5})$

**This is worst-case optimal!**


## Known Lower Bound

Amortized maintenance time: **not**  $\mathcal{O}(|\mathbf{D}|^{0.5-\gamma})$  for any  $\gamma > 0$   
(under reasonable complexity theoretic assumptions)

# IVM $^\varepsilon$ Exhibits a Time-Space Tradeoff

Given  $\varepsilon \in [0, 1]$ , IVM $^\varepsilon$  maintains the triangle count with

- $\mathcal{O}(|\mathbf{D}|^{\max\{\varepsilon, 1-\varepsilon\}})$  amortized time and
- $\mathcal{O}(|\mathbf{D}|^{1+\min\{\varepsilon, 1-\varepsilon\}})$  space.



- Known maintenance approaches are recovered by IVM $^\varepsilon$ .


# Main Ideas in $\text{IVM}^\epsilon$

- Compute the difference like in classical IVM!
- Materialize views like in Higher Order IVM!
- **New ingredient:** Use adaptive processing based on data skew!  
     $\Rightarrow$  Treat *heavy* values differently from *light* values!

# Relation Partitioning

Fix  $\varepsilon \in [0, 1]$  and partition  $R$  into


- a light part  $R_L = \{t \in R \mid |\sigma_{A=t.A}| < |\mathbf{D}|^\varepsilon\}$ ,
- a heavy part  $R_H = R \setminus R_L$ !



# Relation Partitioning

Fix  $\varepsilon \in [0, 1]$  and partition  $R$  into

- a light part  $R_L = \{t \in R \mid |\sigma_{A=t.A}| < |\mathbf{D}|^\varepsilon\}$ ,
- a heavy part  $R_H = R \setminus R_L$ !



Derived Bounds


- for all  $A$ -values  $a$ :  $|\sigma_{A=a}R_L| < |\mathbf{D}|^\varepsilon$
- $|\pi_A R_H| \leq |\mathbf{D}|^{1-\varepsilon}$



# Relation Partitioning

Fix  $\varepsilon \in [0, 1]$  and partition  $R$  into

- a light part  $R_L = \{t \in R \mid |\sigma_{A=t.A}| < |\mathbf{D}|^\varepsilon\}$ ,
- a heavy part  $R_H = R \setminus R_L$ !



Derived Bounds

- for all  $A$ -values  $a$ :  $|\sigma_{A=a}R_L| < |\mathbf{D}|^\varepsilon$
- $|\pi_A R_H| \leq |\mathbf{D}|^{1-\varepsilon}$


Likewise, partition

- $S = S_L \cup S_H$  based on  $B$ , and
- $T = T_L \cup T_H$  based on  $C$ !

# Relation Partitioning

Fix  $\varepsilon \in [0, 1]$  and partition  $R$  into

- a light part  $R_L = \{t \in R \mid |\sigma_{A=t.A}| < |\mathbf{D}|^\varepsilon\}$ ,
- a heavy part  $R_H = R \setminus R_L$ !



Derived Bounds

- for all  $A$ -values  $a$ :  $|\sigma_{A=a}R_L| < |\mathbf{D}|^\varepsilon$
- $|\pi_A R_H| \leq |\mathbf{D}|^{1-\varepsilon}$

Likewise, partition

- $S = S_L \cup S_H$  based on  $B$ , and
- $T = T_L \cup T_H$  based on  $C$ !

$Q$  is the sum of skew-aware views

$$\pi_{()} [R_U(a, b) \bowtie S_V(b, c) \bowtie T_W(c, a)]$$

with  $U, V, W \in \{L, H\}$ .

# Adaptive Maintenance Strategy

Given an update  $\Delta R_*(a, b) = \{(a', b') \mapsto m\}$ , compute the difference for each skew-aware view using different strategies:

Skew-aware View	Evaluation from left to right	Time
$\pi_{()} [R_*(a, b) \bowtie S_L(b, c) \bowtie T_L(c, a)]$	$\Delta R_*(a', b') \cdot \sum_{c'} S_L(b', c') \cdot T_L(c', a')$	$\mathcal{O}( \mathbf{D} ^\epsilon)$

# Adaptive Maintenance Strategy

Given an update  $\Delta R_*(a, b) = \{(a', b') \mapsto m\}$ , compute the difference for each skew-aware view using different strategies:

Skew-aware View	Evaluation from left to right	Time
$\pi_{()} [R_*(a, b) \bowtie S_L(b, c) \bowtie T_L(c, a)]$	$\Delta R_*(a', b') \cdot \sum_{c'} S_L(b', c') \cdot T_L(c', a')$	$\mathcal{O}( \mathbf{D} ^\varepsilon)$
$\pi_{()} [R_*(a, b) \bowtie S_H(b, c) \bowtie T_H(c, a)]$	$\Delta R_*(a', b') \cdot \sum_{c'} T_H(c', a') \cdot S_H(b', c')$	$\mathcal{O}( \mathbf{D} ^{1-\varepsilon})$

# Adaptive Maintenance Strategy

Given an update  $\Delta R_*(a, b) = \{(a', b') \mapsto m\}$ , compute the difference for each skew-aware view using different strategies:

Skew-aware View	Evaluation from left to right	Time
$\pi_{()} [R_*(a, b) \bowtie S_L(b, c) \bowtie T_L(c, a)]$	$\Delta R_*(a', b') \cdot \sum_{c'} S_L(b', c') \cdot T_L(c', a')$	$\mathcal{O}( \mathbf{D} ^\varepsilon)$
$\pi_{()} [R_*(a, b) \bowtie S_H(b, c) \bowtie T_H(c, a)]$	$\Delta R_*(a', b') \cdot \sum_{c'} T_H(c', a') \cdot S_H(b', c')$	$\mathcal{O}( \mathbf{D} ^{1-\varepsilon})$
	$\Delta R_*(a', b') \cdot \sum_{c'} S_L(b', c') \cdot T_H(c', a')$	$\mathcal{O}( \mathbf{D} ^\varepsilon)$
$\pi_{()} [R_*(a, b) \bowtie S_L(b, c) \bowtie T_H(c, a)]$	or	
	$\Delta R_*(a', b') \cdot \sum_{c'} T_H(c', a') \cdot S_L(b', c')$	$\mathcal{O}( \mathbf{D} ^{1-\varepsilon})$

# Adaptive Maintenance Strategy

Given an update  $\Delta R_*(a, b) = \{(a', b') \mapsto m\}$ , compute the difference for each skew-aware view using different strategies:

Skew-aware View	Evaluation from left to right	Time
$\pi_{()} [R_*(a, b) \bowtie S_L(b, c) \bowtie T_L(c, a)]$	$\Delta R_*(a', b') \cdot \sum_{c'} S_L(b', c') \cdot T_L(c', a')$	$\mathcal{O}( \mathbf{D} ^\varepsilon)$
$\pi_{()} [R_*(a, b) \bowtie S_H(b, c) \bowtie T_H(c, a)]$	$\Delta R_*(a', b') \cdot \sum_{c'} T_H(c', a') \cdot S_H(b', c')$	$\mathcal{O}( \mathbf{D} ^{1-\varepsilon})$
	$\Delta R_*(a', b') \cdot \sum_{c'} S_L(b', c') \cdot T_H(c', a')$	$\mathcal{O}( \mathbf{D} ^\varepsilon)$
$\pi_{()} [R_*(a, b) \bowtie S_L(b, c) \bowtie T_H(c, a)]$	or	
	$\Delta R_*(a', b') \cdot \sum_{c'} T_H(c', a') \cdot S_L(b', c')$	$\mathcal{O}( \mathbf{D} ^{1-\varepsilon})$
$\pi_{()} [R_*(a, b) \bowtie S_H(b, c) \bowtie T_L(c, a)]$	$\Delta R_*(a', b') \cdot V_{ST}(b', a')$	$\mathcal{O}(1)$

# Adaptive Maintenance Strategy

Given an update  $\Delta R_*(a, b) = \{(a', b') \mapsto m\}$ , compute the difference for each skew-aware view using different strategies:

Skew-aware View	Evaluation from left to right	Time
$\pi_{()} [R_*(a, b) \bowtie S_L(b, c) \bowtie T_L(c, a)]$	$\Delta R_*(a', b') \cdot \sum_{c'} S_L(b', c') \cdot T_L(c', a')$	$\mathcal{O}( \mathbf{D} ^\varepsilon)$
$\pi_{()} [R_*(a, b) \bowtie S_H(b, c) \bowtie T_H(c, a)]$	$\Delta R_*(a', b') \cdot \sum_{c'} T_H(c', a') \cdot S_H(b', c')$	$\mathcal{O}( \mathbf{D} ^{1-\varepsilon})$
	$\Delta R_*(a', b') \cdot \sum_{c'} S_L(b', c') \cdot T_H(c', a')$	$\mathcal{O}( \mathbf{D} ^\varepsilon)$
$\pi_{()} [R_*(a, b) \bowtie S_L(b, c) \bowtie T_H(c, a)]$	or	
	$\Delta R_*(a', b') \cdot \sum_{c'} T_H(c', a') \cdot S_L(b', c')$	$\mathcal{O}( \mathbf{D} ^{1-\varepsilon})$
$\pi_{()} [R_*(a, b) \bowtie S_H(b, c) \bowtie T_L(c, a)]$	$\Delta R_*(a', b') \cdot V_{ST}(b', a')$	$\mathcal{O}(1)$

Overall update time:  $\mathcal{O}(|\mathbf{D}|^{\max(\varepsilon, 1-\varepsilon)})$

# Materialized Auxiliary Views

$$V_{RS}(a, c) = \pi_{a,c} [R_H(a, b) \bowtie S_L(b, c)]$$

$$V_{ST}(b, a) = \pi_{b,a} [S_H(b, c) \bowtie T_L(c, a)]$$

$$V_{TR}(a, c) = \pi_{a,c} [T_H(c, a) \bowtie R_L(a, b)]$$

- Maintenance of  $V_{RS}(a, c) = \pi_{a,c} [R_H(a, b) \bowtie S_L(b, c)]$

Update	Evaluation from left to right	Time
$\Delta R_H(a, b) = \{(a', b') \mapsto m\}$	$\Delta R_H(a', b') \cdot \sum_{c'} S_L(b', c')$	$\mathcal{O}( \mathbf{D} ^\varepsilon)$
$\Delta S_L(b, c) = \{(b', c') \mapsto m\}$	$\Delta S_L(b', c') \cdot \sum_{a'} R_H(a', b')$	$\mathcal{O}( \mathbf{D} ^{1-\varepsilon})$



# Materialized Auxiliary Views

$$V_{RS}(a, c) = \pi_{a,c} [R_H(a, b) \bowtie S_L(b, c)]$$

$$V_{ST}(b, a) = \pi_{b,a} [S_H(b, c) \bowtie T_L(c, a)]$$

$$V_{TR}(a, c) = \pi_{a,c} [T_H(c, a) \bowtie R_L(a, b)]$$

- Maintenance of  $V_{RS}(a, c) = \pi_{a,c} [R_H(a, b) \bowtie S_L(b, c)]$

Update	Evaluation from left to right	Time
$\Delta R_H(a, b) = \{(a', b') \mapsto m\}$	$\Delta R_H(a', b') \cdot \sum_{c'} S_L(b', c')$	$\mathcal{O}( \mathbf{D} ^\varepsilon)$
$\Delta S_L(b, c) = \{(b', c') \mapsto m\}$	$\Delta S_L(b', c') \cdot \sum_{a'} R_H(a', b')$	$\mathcal{O}( \mathbf{D} ^{1-\varepsilon})$

- Size of  $V_{RS}(a, c) = \pi_{a,c} [R_H(a, b) \bowtie S_L(b, c)]$

$$|V_{RS}(a, c)| \leq |R_H| \cdot \max_{b'} \{|\sigma_{b=b'} S_L(b, c)|\} = \mathcal{O}(|\mathbf{D}|^{1+\varepsilon})$$

$$|V_{RS}(a, c)| \leq |S_L| \cdot \max_{b'} \{|\sigma_{b=b'} R_H(a, b)|\} = \mathcal{O}(|\mathbf{D}|^{1+(1-\varepsilon)})$$

# Materialized Auxiliary Views

$$V_{RS}(a, c) = \pi_{a,c} [R_H(a, b) \bowtie S_L(b, c)]$$

$$V_{ST}(b, a) = \pi_{b,a} [S_H(b, c) \bowtie T_L(c, a)]$$

$$V_{TR}(a, c) = \pi_{a,c} [T_H(c, a) \bowtie R_L(a, b)]$$

- Maintenance of  $V_{RS}(a, c) = \pi_{a,c} [R_H(a, b) \bowtie S_L(b, c)]$

Update	Evaluation from left to right	Time
$\Delta R_H(a, b) = \{(a', b') \mapsto m\}$	$\Delta R_H(a', b') \cdot \sum_{c'} S_L(b', c')$	$\mathcal{O}( \mathbf{D} ^\varepsilon)$
$\Delta S_L(b, c) = \{(b', c') \mapsto m\}$	$\Delta S_L(b', c') \cdot \sum_{a'} R_H(a', b')$	$\mathcal{O}( \mathbf{D} ^{1-\varepsilon})$

- Size of  $V_{RS}(a, c) = \pi_{a,c} [R_H(a, b) \bowtie S_L(b, c)]$

$$|V_{RS}(a, c)| \leq |R_H| \cdot \max_{b'} \{|\sigma_{b=b'} S_L(b, c)|\} = \mathcal{O}(|\mathbf{D}|^{1+\varepsilon})$$

$$|V_{RS}(a, c)| \leq |S_L| \cdot \max_{b'} \{|\sigma_{b=b'} R_H(a, b)|\} = \mathcal{O}(|\mathbf{D}|^{1+(1-\varepsilon)})$$

- Overall: Update Time  $\mathcal{O}(|\mathbf{D}|^{\max\{\varepsilon, 1-\varepsilon\}})$ ; Space  $\mathcal{O}(|\mathbf{D}|^{1+\min\{\varepsilon, 1-\varepsilon\}})$

# Rebalancing Partitions

- Updates can change the frequencies of values in the relation parts!
- This can require rebalancing of partitions.
  - ⇒ Minor rebalancing: Transfer tuples from one to the other part of the same relation!
  - ⇒ Major rebalancing: Recompute partitions and views from scratch!
- Both forms of rebalancing require superlinear time.
- The rebalancing times amortize over sequences of updates.

# Extensions of $\text{IVM}^\epsilon$ ?

## Generalization of $\text{IVM}^\epsilon$

- Partitioning on both attributes of each relation improves space complexity.
- $\text{IVM}^\epsilon$  variants obtain worst-case optimal maintenance time for counting versions of Loomis-Whitney, 4-cycle, and 4-path.

## Ongoing Work

- Characterization of the class of conjunctive count queries that admit  $\mathcal{O}(\mathbf{D}^{0.5})$  worst-case optimal maintenance time
- Implementation of  $\text{IVM}^\epsilon$